The Compromise Effect and Consumer Choice in Repeated Markets

By

Yu-Chun Huang

A thesis submitted in partial fulfilment of the requirements for the degree of
Doctor of Philosophy in Psychology

University of Warwick, Department of Psychology

January 2022

Contents

List of Tables

List of Figures

Acknowledgements

I would like to express my sincere gratitude to Prof. Gordon D. A. Brown for his supportive supervision of my PhD study and tremendous help in finalising this thesis within the limited time frame. His guideline and comments are always invaluable. Also, I am really thankful to his considerable encouragement and solicitude, especially during such a difficult time in Covid-19 pandemics. Having him as my supervisor makes me feel I am the luckiest person in the world.

Moreover, I am very grateful to Dr. Lukasz Walasek for his help for completion of thesis. My thanks also goes to my parents, grandparents, sister, friends, and teddy bear for all the emotional and financial support for me.

Finally, I wish to express my special thanks to the Leverhulme Trust Bridge's Programme and Prof. Thomas T. Hills for financially supporting my PhD.

Declarations

This thesis is submitted to the University of Warwick in support of my application for the degree of Doctor of Philosophy. It has been composed by myself and has not been submitted in any previous application for any degree.

Abstract

The compromise effect in consumer choice is widely observed and has been extensively studied. Standard choice theories typically assume that this context-dependent behaviour will disappear if decision makers have a chance to explore all options to learn their true preferences. If this is true, firms will not be able to exploit consumers at equilibrium. This thesis investigates whether, and under what conditions, the compromise effect may exist in equilibrium in repeated markets even when consumers have extensive market experience. It also studies whether the compromise effect can emerge without attribute trade-offs by focusing on one-dimensional cases. Motivated by relative rank theory and the idea of contextual inference, this thesis develops two psychological models in which consumers' judgements and corresponding choices are influenced both by their "true" preferences and by information that consumers derive from available options in the market. According to relative rank theory, consumers' evaluations of each option are rank-based and constructed from binary ordinal comparisons within the choice set. Using the two different models, this thesis shows that the one-dimensional compromise effect may arise in equilibrium in repeated markets due to consumers' perceptions of their own relative position in the population. Computer simulations of the models show that, under certain conditions, a monopolist may be incentivised to exploit consumers by manipulating the context of market options to take advantage of the compromise effect displayed by consumers. This thesis also reports an experiment to test the basic assumptions of the models. It is concluded that existing models of the compromise effect are inadequate in that they typically fail to allow for the possibility that the compromise effect may persist even in equilibrium when consumers have extensive experience and no attribute trade-off occurs.

Chapter 1: Introduction

This thesis examines how context affects everyday choice. Specifically, it explores conditions under which the context of market alternatives may have permanent effects on consumer choice; that is, influences on the interactions between consumers and a monopolist under market equilibrium[i]. This thesis goes beyond existing models of context-based choice, most of which focus on one-off influences of context on choice and typically fail to address whether the compromise effect (one of the most important context effects)[ii] will persist in the absence of attribute trade-offs[iii] as well as the extent to which the effect may disappear with learning[iv]. Although the models developed in this thesis are intended to address the general issue of the one-dimensional compromise effect in

---

[i] In economics, market equilibrium refers to a situation where the quantity demanded of a commodity is equal to its quantity supplied at a particular price.

[ii] "Context effect" is a psychological term that represents contextual influences on choice. More specifically, it describes the phenomenon whereby a decision between two alternatives can be systematically altered by the presence of other options in the choice set (Prelec, Wernerfelt, & Zettelmeyer, 1997). One of the most widely documented context effects is the compromise effect, which dictates that the choice share of an option can increase when that option becomes the middle option in a consideration set in terms of some physical attribute space (e.g., quality, price) (Simonson, 1989; Tversky & Simonson, 1993). A more detailed introduction of these two effects will be presented in Section 1.1 and 1.2.

[iii] Most of existing models of the compromise effect (e.g., Leong & Hensher, 2014; Sheng, Parker, & Nakamoto, 2005; Simonson, Bettman, Kramer, & Payne, 2013) consider two- or multi-dimensional cases and many of their accounts fail if there is no need for attribute trade-offs. Thus, this thesis goes beyond these models by showing that attribute trade-offs are not necessary for the compromise effect to appear.

[iv] Learning, here and throughout the thesis, means that consumers will infer their own preferences (partially) through experiencing products in the market and will make future purchasing decisions accordingly.

equilibrium, the key issues are first illustrated using consumer choice of food portion sizes as an example.

Over the last four decades, portion and packaging sizes of several nutritionally poor, energy-dense foods have increased (Nielsen & Popkins, 2003; Piernas & Popkin, 2011), potentially leading to increased consumption and hence posing a threat to public health and consumer well-being. The trend toward larger food portions is most manifest and best investigated in the US (Livingstone & Pourshahidi, 2014; Young & Nestle, 2007), but has gradually been mirrored in many other Western countries such as the UK (Wrieden, Gregor, & Barton, 2008), Denmark (Matthiessen, Fagt, Biltoft-Jensen, & Beck, 2003), the Netherlands (Steenhuis, Leeuwis, & Vermeer, 2010), and Australia (Van der Bend, 2017). An associated problem is overeating and increased caloric intake. A number of studies (e.g., Diliberti, Bordi, Conklin, Roe, & Rolls, 2004; English, Lasschuijt, & Keller, 2015; Kling, Roe, Keller, & Rolls, 2016) report that given unchanged perceived levels of biologically determined satiety, expanding portion sizes have a powerful and prolonged effect on amounts of food intake[v]. This phenomenon is known as the portion size effect (for reviews, see Benton, 2015). Moreover, attention has been drawn to the potential contribution of enlarged portion size to the alarming prevalence of overweightness and obesity[vi] (Chandon & Wansink, 2011; Steenhuis & Vermeer, 2009; Young & Nestle,

[v] Several experimental studies on energy-dense foods, including beverages (Flood, Roe, & Rolls, 2006), candies (Marchiori, Waroquier, & Klein, 2011), chips (Vermote, 2018), and a pre-packaged snack (Rolls, Roe, Kral, Meengs, & Wall, 2004) confirm that reducing (increasing) portion size may lead to reduced (increased) caloric intake, with no significant moderating effect of appetite or satiety ratings and food characteristics being found. Moreover, a meta-analysis conducted by Zlatevska, Dubelaar, and Holden (2014) shows that overall consumption rises by 35% with a doubling of portion size.

[vi] For example, BMI is found to be positively correlated with the portion size of some energy dense foods (Rippin, Hutchinson, Jewell, Breda, & Cade, 2019) and of snacks (Albar, Alwan, Evans, & Cade, 2014). Yet, although a positive association between portion size and caloric intake has been identified, there is no established causal link between portion size and

2002), together with criticism of marketing's manipulation of consumer choices and its responsibility in the obesity epidemic[vii].

The portion size effect is just one of many phenomena that imply that consumption choices are driven not just by biological needs but also by contextual cues. The empirical discoveries of this effect have led to a cluster of perplexing questions about contextual influences on judgement and decision making. One prominent research question originates from doubt about the existence of innate, context-independent ideal points. More specifically, if people possess context-independent underlying preferences[viii] that directly inform choices, why and how are observed choices so susceptible to market contexts? This query leads to an important, albeit relatively under-discussed, question regarding the *persistence* of contextual influences on choice. The example of the long-lasting increase in portion size might be taken to suggest that changes in consumer

---

obesity in the literature (Livingstone & Pourshahidi, 2014; Rolls, 2014). This may be because factors contributing to obesity are complex, making it difficult to examine the direct causal relationship between portion size and obesity.

[vii] There is no clear empirical evidence on whether consumers intrinsically prefer larger-sized portions or whether their behaviour are shaped by marketers. However, if consumers demand a supersized product, they could simply buy multiple amounts of it when only small sizes are available. If so, the portion size effect would not be observed as quantity demanded and food consumption would remain constant. On top of that, as pointed out by Nestle (2003), Wansink, and van Ittersum (2007), and Liang, Gemming, Wellard-Cole, and Rangan (2019), recent portion sizes are normally larger than the United States Department of Agriculture's and Australian Dietary Guidelines standard's recommendation. Therefore, it is widely accepted that marketing strategies such as setting very low prices for portion upgrades (Dobson & Gerstner, 2010) and introducing premium loss leaders (Smith & Nagle, 1995) tempt people to consume more than they need and thereby boost the sales of firms (Chandon & Wansink, 2012).

[viii] In the thesis, the term "underlying preferences" is defined as "preferences that people are born with and will not change across choice contexts". This definition is created to distinguish inborn, consistent preferences from other types of preferences (e.g., preferences that are computed during evaluations of options or decision making). A discussion of different types of preferences and the corresponding terms used in the thesis will be provided in Section 1.3.1.

demand for food quantity will endure. In other words, consumers may repeatedly choose and consume suboptimal products even after they have tried the ideal one(s). Perhaps under some contexts, experience *per se* is insufficient to help consumers find their true ideal points and make the corresponding choices, assuming that such ideal points exist. The observations of enduring context effects also imply that there may be other forces than ideal points, such as social norms[ix], that affect consumers' judgements about products and downstream choices as well as contribute to the observation of the long-lasting compromise effect. Yet, if there is no (psychological/ environmental) mechanism that curtails context-dependent choice behaviour, can firms take advantage of context effects by manipulating market context to promote their most profitable products? For example, if selling a larger portion of soft drinks is more profitable, will firms be monetarily incentivised to increases the sizes of soft drinks they produce to the maximum extent that is biologically feasible?

Unfortunately, despite the many efforts that have been made to understand contextual influences on consumer choice, the literature lacks an overarching model that addresses all of these questions simultaneously. Motivated by this lack, the present thesis examines the key issues directly by developing alternative models of the compromise effect based on existing choice theories as well as new empirical evidence. In particular, the thesis will examine the interactive roles of experience and perceived social norms in consumer choice and focus on cases where no attribute trade-off occurs during decision making. Moreover, through investigating shifts of market equilibrium along with context distortion, the thesis explores how firms may make more profit by intentionally taking advantage of consumer bias.

This introductory chapter will begin by reviewing the economic theory of rational choice and its failure to predict behavioural anomalies such as context effects. Then,

[ix] Social influence is crucial in studying consumer choice. To take the example of the portion size effect: The most common explanation is that portion size may signal a social norm and recalibrate people's perceptions of the amount that is appropriate to consume (Herman, Polivy, Pliner, & Vartanian, 2015; Robinson & Kersbergen, 2018). Other research on social influence on decision making will be explored in Section 1.3.4.

several economic and psychological models that account for context-dependent choices will be discussed under three different headings, classified by their assumptions and arguments. The chapter ends with an overview of the novel models that will be proposed in the later chapters.

1.1 Standard economics, rational choice theory, and context effects

"*Economics is the science which studies human behaviour as a relationship between ends and scarce means which have alternative uses*." (Robbins, 1932, p. 15)

In the late eighteenth century, economic studies hinged on a debate concerning the roles of passions, emotions, bodily desires, and reason in human behaviour, mainly based on observing and analysing pricing and purchasing behaviour in a market. The concept of the "rational maximiser" subsequently emerged, but its formal definition and properties were, nevertheless, rather approximate at that time[x]. In the twentieth century, with Robbins' (1932) ground-breaking viewpoint paving the way[xi], rational choice theory invaded the territory of classical economics and soon became a dominant view of standard economic theory. Three crucial assumptions about the nature of rational behaviour became central to mainstream economics.

First, individuals are assumed to possess (or to behave as if they possess) underlying preferences, which are exogenously determined, enduring over at least a certain period,

---

[x] Deriving efficient allocations of resources or optimal actions requires preferences to be well-defined. However, economists used to believe that people act in order to maximise their self-interest or satisfy their personal ends, without formalising the meaning of self-interest and ends, making the theory seem to be virtually tautological (McFadden, 2001).

[xi] Robbins' redefinition extends the scope of economics to a broad investigation of incentive-led behaviour in all types of human activities and social interaction. Being forced to address complex decision-making problems encouraged early economists such as John Hick, Hendrik S. Houthakker, and Paul Samuelson to develop a more rigorous specification of rational behaviour.

stable across choice contexts, and amenable to memory storage and retrieval (Brown, Kingsley, Peterson, Flores, Clarke, & Birjulin, 2008). It is also assumed that preference orderings over choice alternatives are completely context-independent and rank-ordered on an ordinal scale (Sen, 1973, 1977). Further, in order to generate a utility representation [xii] for individuals under perfect certainty, the rationality assumption requires preference orderings to comply with certain formal axioms. Therefore, in mainstream economics, the binary preference relation is normally assumed to be complete, transitive, continuous, strictly monotonic, and strictly convex [xiii] (Hausman, 1992; Houthakker, 1950; Jensen, 1967; Marschak, 1950). Last, individuals are assumed to have an ability to rationally choose the right means to realise a given goal or their own ends. In the language of the standard utility paradigm, this assumption indicates that a rational agent should be able to select the options that maximise utility function subject to

---

[xii] By definition in economics, a utility function is a function mapping consumption bundles or choice alternatives to real numbers. Utility values is transformed subjective values that represent an agent's ordinal preference orderings. Formally, the utility function, u, is said to represent preference if for any bundle $\boldsymbol{x}$ and $\boldsymbol{y}$ and a weak preference relation $\succsim$, $u(\boldsymbol{x}) \geq u(\boldsymbol{y}) \Leftrightarrow \boldsymbol{x} \succsim \boldsymbol{y}$.

[xiii] Let $\boldsymbol{x}$, $\boldsymbol{y}$, and $\boldsymbol{z}$ denote consumption bundles in a feasible set X. Suppose $\succsim$ indicates a weak binary preference relation such that $\boldsymbol{x} \succsim \boldsymbol{y}$ means bundle $\boldsymbol{x}$ is judged at least as good as $\boldsymbol{y}$. The completeness axiom suggests that $\forall \boldsymbol{x}, \boldsymbol{y} \in X, \boldsymbol{x} \succsim \boldsymbol{y} \vee \boldsymbol{y} \succsim \boldsymbol{x}$. The transitivity axiom suggests that $\forall \boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z} \in X, \boldsymbol{x} \succsim \boldsymbol{y} \wedge \boldsymbol{y} \succsim \boldsymbol{z} \Rightarrow \boldsymbol{x} \succsim \boldsymbol{z}$. The continuity axiom suggests that for n = 1,2, …, and $\forall \boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z} \in X, \{\boldsymbol{x}^n\} \rightarrow \boldsymbol{y} \wedge \boldsymbol{x}^n \succsim \boldsymbol{z} \,\forall\, n \Rightarrow \boldsymbol{y} \succsim \boldsymbol{z}$, where $\{\boldsymbol{x}^n\} \equiv (\boldsymbol{x}^1, \boldsymbol{x}^2, \ldots)$ refers to a sequence of options and the arrow → refers to convergence to a limit point. The strict monotonicity suggests that for i = 1,2, …, and $\forall \boldsymbol{x}, \boldsymbol{y} \in X, x_i \geq y_i \forall i \wedge x_i > y_i$ for some $i \Rightarrow \boldsymbol{x} \succ \boldsymbol{y}$. These four axioms are necessary and sufficient for the preference relation $\succsim$ to be represented by at least one continuous utility function. The addition axiom of strict convexity is used to guarantee a quasi-concave utility function, which is necessary for the demand functions generated by utility maximisation to be single-valued, continuous, and differentiable. The strict convexity axiom suggests that $\forall \boldsymbol{x}, \boldsymbol{y} \in X, \boldsymbol{x} \succsim \boldsymbol{y}, \boldsymbol{y} \succsim \boldsymbol{z} \wedge \boldsymbol{x} \neq \boldsymbol{y} \Rightarrow \lambda \boldsymbol{x} + (1-\lambda)\boldsymbol{y} \succ \boldsymbol{z} \,\forall\, 0 \leq \lambda \leq 1$.

prevailing constraints (Afriat, 1967). Altogether, standard economics generally describes decision makers as being optimising, deliberative, autonomous, and self-regarding, in the sense that they always behave as if they meticulously evaluate all the possible outcomes and choose the alternative with the highest utility value based on their innate, static underlying preferences.

While the approach that standard economic theory employs to approximate human decision-making has occupied a dominant position in social science for nearly a half-century, it has long been criticised for its inability to offer a coherent and rigorous account of observed behaviour as well as lack of predictive capability under some circumstances (Friedman, Isaac, James, & Sunder, 2014; Levin, Gaeth, Schreiber, & Lauriola, 2002; Maniadis, Tufano, & List, 2014; Tversky, 1972). Many of these criticisms revolve around the independence of irrelevant alternatives (IIA) and the regularity principle – two important properties of rational choice that serve as a root of many economic models. In accordance with Luce (1959), the former states that the relative choice probabilities of two alternatives are independent of the choice set in which the alternatives are presented[xiv]. By extension, the IIA posits that the introduction of a new option to the available choice set must not change the preference order of any pre-existing alternatives[xv] (Bettman, Luce, & Payne, 1998). Further, the IIA connotes the regularity principle, which states that the probability of choosing, or the absolute market share of, an option from a choice set cannot be increased by enlarging the set (Rieskamp, Busemeyer, & Mellers, 2006). Overall, these two principles imply that if decision makers satisfy the aforementioned assumptions of rationality, their choices should not, in theory, be influenced by the presence of other options.

[xiv] At the time of the IIA's inception, its definition often varies with authors (e.g., Arrow, 1951; Luce, 1959; Radner & Marschak, 1954). Although those definitions were claimed to be essentially the same by their respective author, they are in fact quite different (Ray, 1973). To avoid confusions, this thesis will use the definition developed by Luce (1959).

[xv] This extended definition presupposes that choice is purely and directly governed by a well-defined underlying preference.

However, in practice, numerous counterexamples to the IIA and regularity principles have been observed in both human (e.g., Tversky & Simonson, 1993; Chernev, 2004) and non-human subjects (e.g., Latty & Beekman, 2011; Lea & Ryan, 2015; Shafir, Waite, & Smith, 2002). Indeed, shortly after Luce (1959) positioned the IIA as a choice axiom, Becker, DeGroot, and Marschak (1963) reported a violation of the IIA. As demonstrated in their work, adding an option to a choice set takes disproportionately more share of choices from similar ones than from dissimilar ones. This was later termed the "similarity effect" by Tversky (1972). A decade later, Luce (1977) declared that regularity was the only axiom of general choice probabilities that had not been experimentally violated. Yet, history repeats itself. In 1982, Huber, Payne, and Puto used a simple laboratory study to confirm that the mere addition of asymmetrically dominated or relatively inferior alternatives increased the absolute share of the alternative that dominates it. This is known as the attraction effect. Henceforth, research on this topic continues to flourish, with a growing body of evidence demonstrating the fact that choice behaviour is often affected by the composition of the choice context, which runs counter to the traditional assumption that options are ranked independently. Together with more recent demonstrations of other phenomena, such as the compromise effect[xvi] (Simonson, 1989; Tversky & Simonson, 1993), and the attribute-balance effect[xvii] (Chernev, 2004), the context-relevant behavioural anomalies are grouped together as context effects.

[xvi] The compromise effect was first identified by Simonson (1989), who distinguished it from another two context effects – the attraction effect and the similarity effect. Although it is distinct from those two effects, its violation of regularity (Kivetz, Netzer, & Srinivasan, 2004; Simonson, 1989) and the IIA (Hutchinson, Kamakura, & Lynch, 2000) are found in many experiments.

[xvii] The attribute-balance effect suggests that a choice alternative with equal attribute values will be perceived as a compromise option even when it is not positioned in the middle of the choice set (Chernev, 2004).

1.2 The compromise effect

To answer the research question outlined at the beginning, the present thesis concentrates on the causes and consequences of one of the most-studied context effects, namely the compromise effect, and its effect on a firm's design of product line, market equilibrium points, and consumer welfare.

By definition, the compromise effect occurs when the choice share of an option is boosted when it becomes the compromise (i.e., middle) option as a new, non-dominating option is introduced to the original choice set (Simonson, 1989). The compromise effect implies that choices are sensitive to the relative positions of options in a consideration set, with options that are positioned between other non-dominating options being chosen more frequently than extreme ones (Müller, Vogt, & Kroll, 2012). While the compromise effect contrasts markedly with standard economic theory, it has been documented across various domains and product categories[xviii], ranging from the demand for wine (McFadden, 1999) and soft drink size choices (Sharpe, Staelin, & Huber, 2008) to healthcare-related decisions (de Bekker-Grob & Chorus, 2013) and choices of portable computers and speakers (Kivetz, Netzer, & Srinivasan, 2004). In addition, its practical implications have been found useful in altering dietary intake (Carroll & Vallen, 2014; Pinger, Ruhmer-Krell, & Schumacher, 2016; Sharpe, Staelin, & Huber, 2008; Wu, Gong, Chen, & Hu, 2020).

Technically, there are two basic forms of choice-set type employed to examine the compromise effect. In most studies, the change of choice shares is compared under either a binary and a trinary choice set or using two triplets, with the binary-trinary set comparisons receiving more attention. The novel models proposed in the present thesis will use equal-sized set comparisons (e.g., in a quintuplet-quintuplet case) to investigate

---

[xviii] According to a meta-analysis conducted by Neumann, Böckenholt, and Sinha (2016), the average magnitude of the compromise effect exhibits substantial variation across the uses of product and attribute types (e.g., durable vs. nondurable goods, hedonic vs. utilitarian products, and numeric vs. non-numeric attributes) and experimental methodologies (e.g., using binary-trinary choice sets vs. trinary-trinary choice set, numbers of trade-off dimensions, etc.).

the psychological mechanism behind the one-dimensional compromise effect, together with its influence on a firm's optimal menu, market equilibrium points, and consumer well-being.

1.3 Models of the compromise effect

The existence of context effects undermines some assumptions of rationality, with some arguing that it proves that humans are irrational, but others defending the idea that the definition of rationality in standard economics is too narrow and that departing from this narrowest sense of rationality cannot be counted as irrational behaviour (Thaler, 1992). Nonetheless, it is undeniable that the violation of the IIA and regularity reveal that anomalous behaviour belies at least one assumption of rational choice and is incompatible with many choice theories. In response to this crisis, economists and psychologists have developed many theoretical explanations built upon various modified assumptions and experimental evidence.

Most behavioural economic models of context effects retain the classical utility paradigm with additional ad hoc modifications such as adding an extra element like the relative position of options (Sharpe, Staelin, & Huber, 2008) or an uncertainty parameter (Guo, 2016) or introducing psychological factors such as salience (Bordalo, Gennaioli, & Shleifer, 2013) or limited attention (Manzini & Mariotti, 2014) to a utility function. By contrast, psychological models are less likely to assume the existence of innate, context-independent underlying preferences defined over attributes. Far fewer are inclined to involve utility calculation, with some even assuming no value computation (Vlaev, Chater, Stewart, & Brown, 2011). Nevertheless, several psychological models (e.g., Howes, Warren, Farmer, El-Deredy, & Lewis, 2016; Ronayne & Brown, 2017; Wernerfelt, 1995) implicitly hold that people make choices in order to maximise their utility from consumption, which is in line with the assumption of maximisation.

The rest of this section will explore some important choice theories and models that may help to explain the compromise effect. Section 1.3.1 will briefly introduce the debate about preference stability, which is vital in understanding the theoretical implications of context effects and assumptions of relevant models. The subsequent subsections will review prior economic and psychological models of the compromise effect.

1.3.1 Preference stability

One intuitive interpretation of context-dependent choices is that an individual's underlying preference over a choice alternative or an attribute is influenced by the environment in which the decision is made (Grether & Plott, 1979; Sen, 1993; Simonson & Tversky, 1992) since choice is typically thought to be purely achieved by preference maximisation. This viewpoint not only questions the assumption of preference stability, but limits the applicability of the revealed preference approach, a conventional economic method for inferring people's preference from observed behaviour under different circumstances.

The idea of revealed preferences was first introduced by Samuelson (1938), at a time when economists were eager to abandon psychological foundations and aspired to positivism. Unlike the approach of psychologists who strive to uncover the black box of a decision-making process, the theory of revealed preference works backwards by assuming that a set of preferences that governs choices can be empirically deduced from analysing observed behaviour, without needing to assume any functional form of utility (Richter, 1966) or any underlying psychological or biological causes such as moral standards, personality, patience, cleverness, etc. (Binmore, 2008). More specifically, this theory entails that choosing a certain consumption bundle over others reveals that the chosen one is preferred to other affordable consumption bundles in the consideration set, given constant income and prices (Samuelson, 1948). Hence, a representative model of revealed preferences can be legitimately inferred by altering circumstances such as income or prices or both. Yet, under this framework, context-sensitive choices may produce revealed preference orderings that are unstable, or may even reverse, across contexts, which in turn incapacitates the revealed preference approach[xix].

[xix] The weak axiom of revealed preference argues that if one product is purchased rather than another, then it must be so in every case unless there is a change in relevant information for decision-making such as prices, incomes, product qualities, and so on. In straightforward terms, it is assumed that choice is consistent under the same circumstance (Mariotti, 2008). This axiom is required because revealed preference orderings are generated by recording a decision of each

The evidence of seemingly labile underlying preferences and the failure of the revealed preference approach accelerated the development of an extreme perspective according to which there are no innate, context-independent underlying preferences that can be stored, retrieved and revealed (Lichtenstein & Slovic, 2006; Payne, Bettman, & Johnson, 1992; Weber & Johnson, 2009), and that therefore preferences are typically constructed at the time of decision, probably through using arbitrary cues as anchors (Bettman, Luce, & Payne, 1998; Slovic, 1995; Stewart, Chater, & Brown, 2006). However, as I will argue in this thesis, context effects may be mistakenly construed as an evidence of preference construction. As pointed out by Warren, McGraw and van Boven (2011), the existing literature usually defines preference construction in terms of computation but operationalises it as context sensitivity, since researchers often confuse preference incompleteness with preference instability. More specifically, the need to construct preferences at the time of decision stems from the lack of a complete underlying preference and/or incompetence in discerning a pre-existing preference, rather than from preference instability. Put differently, if people are endowed with an underlying preference that is always retrievable but context-dependent, whereby one is preferred to another in one context but the reverse in another context, then people would not need to calculate a specific preference for the choice set (thanks to retrievability) but would still exhibit context-dependent behaviour. In short, preference construction is not inherently incongruent with preference stability, nor does stability necessarily imply complete innate preferences (Bettman, Luce, & Payne, 2008). Accordingly, the first step in appropriately interpreting the implications of context effects for preference stability is to investigate whether there exists a set of innate underlying preferences for each attribute or product, followed by examining their properties such as stability, retrievability, and completeness.

Due to technological and methodological issues, definitive evidence as to the existence of an innate underlying preference has not yet been gathered. Nor has research

---

pairwise comparison within the set (Rulli & Worsnip, 2016). If observed choices reverse when a context changes, or less abstractly, if IIA is violated, the revealed preferences, inferred from pairwise comparisons, may be inconsistent with the choice made in the original set.

hitherto directly examined the properties of such preferences if they do exist. In behaviour genetics, robust heritable effects on variation in attitudes and behaviours, including overconfidence (Cesarini, Lichtenstein, Johannesson, & Wallace, 2009), prosocial tendency (Hur & Rushton, 2007), political orientations (Funk et al., 2013), party affiliation (Settle, Dawes, & Fowler, 2009), and investment decisions (Cronqvist & Siegel, 2014), have been identified. Moreover, Simonson and Sela (2011) report that the propensity to choose the compromise option has a heritable component. They argue that this genetic predisposition reflects the fact that people are born with different degrees of susceptibility to behaving in a way that departs from standard economic predictions. Although these studies appear to confirm that human dispositions and individual differences have a significant genetic component, they are not sufficient to substantiate genetically determined underlying preferences. This is because such evidence cannot exclude the possibility that what genes really affect is the process of preference formulation. For example, variations in the level of enjoying the burning sensation of chili peppers among people who grew up in the same village (Rozin & Schiller, 1980) may come from individuals' inborn differences in the ability to learn to eat foods containing chili peppers. Overall, then, past evidence on innate underlying preferences is still fragmentary.

Even if innate preferences exist, it is difficult to empirically prove whether they are context-dependent or not. Indeed, acknowledging context effects as cogent evidence of preference instability is only "in a manner of speaking" since the link between choice and underlying preferences is more intricate than what standard economic theory, including the revealed preference approach, assumes. More concretely, as there are numerous factors that interact with innate underlying preferences in affecting decision-making, context sensitivity may result from other causes or their joint effects. For instance, recognition of each option's objective information, amounts of experiences in selecting under the same choice set, and a tendency to retrieve underlying preferences may be time- and context-dependent (Natenzon, 2019; Warren, McGraw & van Boven, 2011). Thus, a revealed preference elicited from a group of observed choices can reflect multiple forces under each choice context, rather than a simple innate preference (Smith, 2008). Observed choice behaviour therefore does not necessarily imply properties of an innate preference.

Behavioural scientists typically admit the existence of context-insensitive preferences that are genetically or evolutionarily formed and hard-wired (Durante & Griskevicius, 2016). One reason behind this is the limited effects of environment on stated preferences and choices. A simple example could be a preference for a warm over a cold or hot ambient temperature (Hsee, Yang, & Shen, 2009). Evidently, neither living environments nor choice contexts can cause people to prefer exposure to extreme temperatures. This logic can be applied to preferences in other domains such as instinctive tastes for food attributes. Furthermore, as suggested by Simonson (2008), although the assumption of innate stable preferences is still non-falsifiable to date, it is a convenient and powerful device that facilitates theorisation of observed behaviour phenomena and generation of novel insights.

Most existing models of the compromise effect are based on the default assumption of preference stability, with some involving preference construction. The definition of "preference" is slightly slippery in the literature. To clarify, throughout this thesis, the term "underlying preference" will be used to denote a biologically determined preference, that exists before the preference objects are considered. In contrast, the term "inferred preference" indicates a preference that individuals learn[xx], form, or perceive during the process of choosing. Conceptually, inferred preferences more or less resemble constructed preferences. In addition, to capture the possibility that people may not choose solely in light of their inferred preferences, the term "expressed preference" is employed to denote observed preference orderings. This classification of preferences allows a more nuanced investigation into the causes of anomalous behaviour, than current literature does.

1.3.2 Group 1: Preference uncertainty

The first group of accounts of the compromise effect consists of a triad of classic explanations, namely, justification (Simonson, 1989), the extremeness aversion principle

---

[xx] It is important to note that in the models proposed by the thesis, inferred preferences will (at least partially) come from experience consuming products. "Learning" will thus require physical consumption.

(Simonson & Tversky, 1992), and expected-loss minimisation (Sheng, Parker, & Nakamoto, 2005). Rather than arguing about the nature of preferences, these accounts assign considerable importance to the issues of uncertainty and decision conflicts, positing that in conditions of preference uncertainty, people often have difficulty in making a choice, especially when the decision task involves tackling a conflict associated with how much of one attribute to forgo in favour of another. In order to cope with this problem, it is asserted that people tend to resort to pre-coded strategies depending on the choice environment, and consequently may behave as if they systematically violate rationality.

The need to justify one's choice is the first documented explanation of the compromise effect. This narrative account was conceived by Simonson (1989), suggesting that compromise behaviour arises from a search for reasons either to oneself or to others as well as a desire to be positively appraised by others. More precisely, compromise/ middle options are usually perceived as safer and more sensible options as they can be easily justified by stating that they contain more balanced attribute values, which potentially helps reduce the likelihood of being criticised (Dhar & Simonson, 2003; Simonson, 1989). This point of view implies that people often expect their choices to be assessed by others. It is one of the few explanations of context effects that mention the importance of interpersonal influence but, unfortunately, it is not further formalised as a theory. On top of that, the notion of justification is apparently fashioned *ad hoc* and oversimplifies the motivations behind decision making.

Some years later, Itamar Simonson and Amos Tversky proposed a more thorough account of the compromise effect, known as the extremeness aversion principle. Extending the theory of loss aversion[xxi] to include relative advantages and disadvantages, Simonson and Tversky (1992) argue that individuals tend to compare advantages and disadvantages of options with each other with respect to every attribute dimension.

[xxi] Loss aversion was first introduced by Kahneman and Tversky (1979), and is arguably the most prominent aspect of prospect theory. The concept of loss aversion suggests that individuals tend to overweight real or potential losses relative to equivalent gains (for reviews, see Sokol-Hessner & Rutledge, 2019).

Usually, options have relative advantages in some dimensions and disadvantages in others and intermediate options have small disadvantages in relation to the other options. As, by the definition of loss aversion, disadvantages are weighted more heavily than advantages and marginal values of attributes is diminishing, middle options are normally perceived to be more attractive than other options. This implies that if more than one attribute needs to be evaluated, the possibility of attribute conflict may increase, which would enhance preference uncertainty and then drive extremeness aversion. Put the other way round, this account tacitly assumes that the compromise effect may be hard to observe in the absence of attribute trade-offs, making it less useful in explaining one-dimensional cases such as a change in portion sizes.

Likewise, the idea of expected-loss minimisation (Sheng, Parker, & Nakamoto, 2005) suggests that uncertainty motivates people to undertake value computations in a decision-making process. More specifically, their model is built based on the idea that uncertainty affects choice through perceived risk of choice outcome. To capture the effect of risk on choice, Sheng, Parker, and Nakamoto (2005) use expected loss as a proxy for measuring risk. Combined with the fact that minimising expected loss is equivalent to maximising expected gain, when there is uncertainty as to which choice alternative can produce the highest value, Sheng, Parker, and Nakamoto (2005) argue that consumers tend to minimise expected losses during decision making in the presence of uncertainty. Mathematically, the expected loss of a potential choice is calculated as a summation of the probability of an alternative i being the best choice multiplied by the difference between subjective values of the potential choice and the alternative i. Accordingly, individuals who encounter greater uncertainty about subjective values, as reflected by a more uniformly distributed probability of an option being the optimal one, are more likely to believe middle options are the best options and therefore exhibit compromise behaviour. From this perspective, attribute conflicts may increase the likelihood of the compromise effect through increasing preference uncertainty.

In summary, none of these three accounts specify the properties of preferences. Instead, the main argument is that when people have little information about preferences or subjective values in advance, they are more likely to use strategies to make an optimal decision. This attempt however may make their behaviour seem biased. On the positive side, this implies that once preference uncertainty is resolved, choice will be consistent with the true (underlying) preference, no matter what it is, as people are maximisers[xxii]. To date, this viewpoint has been indirectly supported by numerous empirical studies. For example, experimental results show that middle options are less likely to be chosen when decision makers have complete information about products (Chuang, Kao, Cheng, & Chou, 2012), when they have more knowledge concerning attribute or option differences (Sheng, Parker, & Nakamoto, 2005), and when their private self-awareness or self-confidence in decision making is greater (Chuang, Cheng, Chang, & Chiang, 2013). Since preference uncertainty is generally assumed to arise from insufficient information or a lack of knowledge about options, these findings partly support the idea that less uncertainty makes people more likely to select what they would really like rather than choosing the compromise one.

1.3.3 Group 2: Informative menu

Along the lines of the first group, models in the second group also recognise the role of uncertainty in the compromise effect, while differing in explicitly assuming an accessible and stable underlying preference and in adopting the standard utility paradigm to model individual behaviour. More importantly, this group of models considers context-dependence that naturally arises through rational inference based on market context during decision-making processes. Three behavioural economic models that represent three distinct facets of this account are described below.

Firstly, it is argued in some economic literature that context effects may be endogenously driven by context-dependent beliefs. Piermont's (2017) model provides an

---

[xxii] This implication does not mean the reason-based justification and extreme aversion accounts exist to resolve the uncertainty. What it actually means is that theories addressed in this subsection attribute the compromise effect to uncertainty and believe that people deliberately adopt various strategies to make an optimal decision. Therefore, it is expected that when uncertainty is resolved (by whatever reasons), these "maximisers" will choose strictly according to their underlying preferences.

example. In this model the compromise effect is characterised as a shift in the chooser's beliefs about the state space, which is *ex ante* uncertain and can be inferred through treating a choice set as a signal about the possibility of different states [xxiii]. More specifically, since the utility of an option is contingent on which state of the world is realised, choosers have to maximise expected utility with respect to the menu-induced belief over the state space, making their choices likely to change with contexts in response to the updated beliefs. To illustrate the idea of the model less abstractly, Piermont (2017) provides the following example of equilibrium market behaviour. Consider a situation where firms in the market are privately endowed with a profile associated with the quality of products they supply and the cost of production. Suppose that firms differentiate themselves by setting different menus, and that consumers are perfectly aware of this signalling mechanism. To differentiate from low-type firms, high-type firms are inclined to introduce additional high-quality, high-cost products, which are expensive for low-type firms to carry. Therefore, in equilibrium, consumers can infer firms' types by observing menus. As proved by Piermont (2017), consumers' optimal decisions under contexts provided by different types of firms are different. The compromise effect occurs because consumers choose low-quality, low-priced options (e.g., the first option in an ordered choice set) or an outside option when they believe the menus are provided by low-type firms, while choosing slightly better but more costly options (e.g., middle options) when menus are thought to be offered by high-type firms due to an inclusion of additional high-

---

[xxiii] Piermont's (2017) model is an application of signalling theory in economics, which was initially developed by Spence (2002) with the aim of showing how two parties (e.g., prospective employees and organisations) engage in activities that may lower the level of asymmetric information between them. According to Spence (2002), when there is information asymmetry, one party, who holds information about the true state that is hard to be observed by others, will choose a means to credibly signal the state (e.g., use an education qualification to signal one's competence), and the other party will choose a way to interpret the signal. In Piermont's model, the monopolist will treat a menu as a signal of quality of its products and consumers will interpret the menu in a sensible way to obtain information about products' quality.

end products. Conversely, if consumers know with certainty which state of the world was to be realised, choices would be consistent across menus as consumers' underlying preferences for alternatives, and thus utility and *ex post* tastes, are assumed to be the same across menus.

Interestingly, it is postulated that in a world without state uncertainty, the compromise effect may still arise when there is a sufficient proportion of uninformed consumers who infer their preferences over available products from payoff-relevant information conveyed in menus. Wernerfelt (1995) suggests that knowledge of one's own relative position in the population distribution of tastes is enough for the optimal product to be identified if the entire menu is observed[xxiv]. Following from this theory, Kamenica (2008) argues that even though all consumers correctly know their relative, but not absolute, tastes and recognise the overall distribution of tastes, the distinction between technical (e.g., lumen) and hedonic (e.g., brightness) units of quality often makes a proportion of consumers unsure about the hedonic interpretation of quality. Therefore, when products are described in technical units, these uninformed consumers need to rationally infer their preferences from market offerings, which may consequently trigger seemingly irrational behaviours. More concretely, these uninformed, rational agents can be assumed to understand that an introduction of high-quality, high-priced products signals that there exists a great number of consumers who have a high hedonic value of technical units of quality (e.g., brightness per lumen) since it is not profitable to introduce the options if consumer characteristics are otherwise. As proved by Kamenica (2008), uninformed agents may rationally speculate that they are more likely to have high willingness to pay per technical unit of quality when a firm offers multiple products than when just one product is sold in the market. Deductively, the overall demand by uninformed consumers for the second option is greater in a trinary set relative to that in

---

[xxiv] This is based on the viewpoint that product lines are designed to capture consumers' tastes as much as possible. Hence, rank among products is believed to correspond to a consumer's rank in true preference. This idea and relevant experimental evidence will be explored in detail in the next subsection.

the binary set – a quintessential exhibition of the compromise effect. Furthermore, Kamenica (2008) underlines the fact that if many consumers rely on context to make an inference, the monopolist will be incentivised to manipulate the informational content of the menu by introducing a costly option of overly high quality that is unprofitable per se, but effectively improves the demand for other options. This indicates that supply distortions and compromise behaviour may be present in market equilibrium as long as there exists a sufficient number of uninformed consumers.

Offering an additional perspective, Guo (2016) develops a novel decision-making theory called contextual deliberation, which approaches the idea that a choice set is informative from a completely different angle. As argued by Guo (2016), market context systematically affects the purchase decision purely through its influence on pre-choice deliberation. This approach therefore differs from the above-mentioned contextual inference models in which context *per se* carries payoff-relevant information. In Guo's (2016) model, underlying preferences about available products are *ex ante* uncertain, even when product attributes are known. This uncertainty, however, can be (partially) resolved by engaging in costly deliberation such as retrieval or revelation of preference cues from memory, reflection of personal need, inspection of product specification and so on. Since introduction of a new choice alternative to the set may increase the benefits of garnering information, the incentive for deliberation is enhanced[xxv]. Thereby, choice probability will be increasingly reallocated from prior preferred alternatives that are usually chosen by feeling to other alternatives that are *ex ante* not preferred but are optimal under

[xxv] This more or less corresponds to the viewpoint of the first group, whereby difficulty of choice tasks may dispose people to deliberately use cognition to make a decision, and thus exhibit the compromise behaviour. However, Guo (2016) also notes that it is not always the case that expanding the choice array can boost the incentive to deliberate. For instance, the level of deliberation is expected to decline if the newly added alternative obviously dominates other pre-existing options.

deliberation; that is, the compromise options[xxvi]. Yet, it is important to note that the compromise effect may not emerge when the cost of deliberation exceeds a reasonable range, because this may substantially reduce the incentive to deliberate when the choice set expands (Guo, 2016).

Despite large differences in proposed decision-making mechanisms, these three models share similar assumptions. First, models in this group attempt to explain seemingly irrational behaviour like the compromise effect while embracing traditional assumptions of rationality used in mainstream economics, including the general axioms of preferences and the expected utility paradigm. More specifically, these models are built on the assumption that individuals are, to some degrees, aware of their underlying preferences regarding existing products[xxvii] and may rationally maximise happiness gained from consumption. However, in respect of the potential effects of experiences on decision making, these models have slightly different viewpoints. Though it is not explicitly mentioned in Kamenica's (2008) paper, it appears that uncertainty concerning the value transformation from technical units to hedonic units can be reduced by experience. It is then reasonable to predict that the compromise behaviour may disappear

---

[xxvi] The notion that the compromise options are more likely to arise from deliberate information processing, instead of the adoption of some choice heuristics or simply being based on feeling has much empirical support (Dhar & Gorlin, 2013). The compromise effect is found to be less visible if participants' cognitive resources are depleted, for example, by a pharmaceutical reduction of brain serotonin levels (Lichters, Brunnlieb, Nave, Sarstedt, & Vogt, 2016) or by increased time pressure (Pettibone, 2012).

[xxvii] Piermont (2017) clearly points out that underlying preferences for *ex-post* outcomes are fixed and consumers are aware of their state-dependent preferences of products. In other words, consumers know their preference orderings over products in each state. Slightly differently, Guo (2016) assumes that although people realise their underlying preferences, the memory retrieval of these preferences in a decision task may be structurally affected by the characteristics of the decision context. As for Kamenica (2008), consumers are assumed to only know their relative position in the population distribution of tastes, although they are endowed with a stable underlying preference.

in conditions of repeated purchasing. Likewise, albeit in a slightly different vein, Piermont (2017) surmises that rational agents will adjust their perceptions of the information conveyed in a given menu based on their experiences. This makes people's beliefs about the state space aligned to the true probability of states and reduces the probability of manipulation by the supply side. On the other hand, Guo's (2016) analysis implies that experience or knowledge of products may increase the compromise effect through lowering the cost and/ or boosting the need of deliberation. The role of experience in the relationship between uncertainty and the compromise effect will be explored in more detail in a later subsection.

1.3.4 Group 3: Social norm models

As mentioned, Kamenica's (2008) explanation of the compromise effect relies on the assumption of a rank-order decision rule [xxviii] proposed by Wernerfelt (1995). According to this theory, decision makers normally believe that the choice array reflects the distribution of absolute tastes or valuations in the population. Therefore, in the presence of uncertainty about preferences and/ or unavailability of attribute information, decision makers may leverage comparative information about themselves and available options to infer the optimal decision. One possible way of doing this is to estimate their relative position among others on a relevant taste, and then choose the alternative whose ranked position in the market matches decision makers' estimated relative standing in the population (Burson, 2007; Simonson, 1993; Wernerfelt, 1995). If the market context precisely reflects the distribution of taste and the estimate of relative taste is accurate, choices made using the rank-order decision rule will be identical to the utility-maximising choice (Wernerfelt, 1995).

However, this choice rule also implies that a shift of choice sets can profoundly alter decision makers' inferred preferences and in turn affect the resulting choice through

---

[xxviii] This choice rule is also referred to as the matching strategy by Burson (2007). The term "matching" here refers to the process of selecting an option by aligning its rank in an ordered choice set with the decision maker's perceived rank in a taste distribution.

altering perceived social norms, while leaving the estimate of own relative position in the population unchanged. Based on this idea, Wernerfelt's (1995) theory explicates a possible mechanism underlying context-dependent behaviour. Indeed, Wernerfelt (1995) uses two examples to formally illustrate that the rank-order decision rule may ultimately give rise to the compromise effect if decision makers first face a smaller and then an extended set and are unsure about the population rank of products provided in the initial set since they know the set is censored. Importantly, in both examples, consumers are assumed to know their true relative taste in the population. The compromise effect is ascribed to uncertainty about interpreting preference-relevant information revealed by censored choice sets that first appear because consumers do not know on which side the censoring occurs.

This theory is supported by multiple experimental studies. Firstly, a preliminary survey briefly mentioned in Prelec, Wernerfelt, and Zettelmeyer (1997) shows that short participants often selected the shortest waterproof poncho in the choice set, even though with perfect information, the optimal one for almost all participants would be the largest size. Furthermore, in a more thorough experiment, Prelec, Wernerfelt, and Zettelmeyer (1997) find that across eight distinct product categories, including air conditioners, cameras, coffeemakers, etc., participants do choose the product whose relative position in the set on the price-quality Hotelling line matches their estimated relative willingness to pay in the population distribution. More importantly, their result indicates that, on average, this inference effect alone explains two-thirds of the magnitude of the compromise effect. This result constitutes powerful evidence for Wernerfelt's (1995) theory.

The idea that rational inferences shape choices is confirmed empirically by other researchers such as Burson (2007), Gershoff and Burson (2011), and Hamilton, Ratner, and Thompson (2010). Yet, in contrast to the theoretical assumptions of Wernerfelt (1995) and Kamenica (2008), many of these studies demonstrate that individuals' self-perception about their relative standing in population is usually distorted. Bias arises since choice is made by aligning options with this inaccurate self-perception. One possible, and straightforward, reason behind the erroneous self-perception is people's poor judgement of their own absolute standing on dimensions such as taste, ability, attitude, belief, and behaviour. Take the example of ability. A low correlation between self-assessed and objectively measured ability has often been reported in past research (e.g., Freund & Kasten, 2012; Krueger & Mueller, 2002; Zell & Krizan, 2014). More importantly in the present context, several experimental results (e.g., Burson, 2007; Burson, Larrick, & Klayman, 2006; Kruger, 1999) find that people's assessments about their skills relative to others can be manipulated by assigning them to either a hard or easy task. Furthermore, as pointed out by Burson (2007), manipulation of task difficulty that significantly alters an individual's perceived comparative ability may consequently lead to downstream behaviour of inconsistent product choices. This finding suggests that suboptimal choice is likely to stem from imprecise perception of comparative position, rather than arbitrary changes in sizes of choice sets.

Incorrect estimation of the population distribution is another crucial factor responsible for the flawed estimate of relative position. In the psychology literature, this type of judgemental error is frequently attributed to the tendency to overestimate the variability or dispersion of the population distribution (Gershoff & Burson, 2011; Judd, Ryan, & Park, 1991), which can result from a memory bias for extreme examples (Gershoff, Mukherjee, & Mukhopadhyay, 2007; Madan, Ludvig, & Spetch, 2014; Nisbett & Kunda, 1985). More specifically, even when people are fully aware of their true absolute standing, their overestimation of the popularity of extremes may generate an estimate of overly dispersed distribution of population, which leads to biases in estimating relative position (Burson & Gershoff, 2015). Accordingly, Gershoff and Burson (2011) argue that simply presenting extreme examples of others to a person can effectively shift his/ her perceived relative position in population and observed choices. Based on this perspective, it is predicted that better remembering extremes may lead people to feel their relative position is close to the middle, resulting in compromise behaviour when the rank-order decision rule is adopted.

A separate line of scholarship holds instead that incorrect estimation of population distribution may result from extrapolating unrepresentative samples obtained from social circles to the population. This stream of research (e.g., Galesic, Olsson, & Rieskamp, 2012, 2018; Hertwig, Pachur, & Kurzenhäuser, 2005; Juslin, Winman, & Hansson, 2007) shows that a majority of people seem to be capable of accurately perceiving information in their social contacts, but the direct use of the social-circle information to infer population properties[xxix] may predispose them to suffer prediction errors. It is likely that the errors stem from the fact that in many situations people prefer to interact with like-minded others and move in similar social circles (Currarini & Mengel, 2016; McPherson, Smith-Lovin, & Cook, 2001; Smirnov & Thurner, 2017), inevitably yielding far lower variances in individuals' social circles, relative to the real population variances, as well as to unrepresentative samples being generated (Galesic, Olsson, & Rieskamp, 2018).

Not surprisingly, group homophily, together with its contribution to biased population estimates, has been considered a possible key driver of the false consensus effect (Kitts, 2003), a well-documented phenomenon referring to an overestimation of prevalence of one's own views, traits, and preferences in a general population[xxx] (Marks & Miller, 1987; Mullen et al., 1985). As reported by Galesic, Olsson, and Rieskamp (2018) and Bruine de Bruin, Galesic, Parker, and Vardavas (2020), the size of the false consensus effect is positively associated with the extent of homophily in individual social circles. Extending this notion, it is speculated that people who encounter strong in-group

---

[xxix] These studies show that regardless of representativeness, people like using their immediate social circles to estimate broader social environments, while population estimates are found to be smoother than social-circle distributions. According to Galesic, Olsson, and Rieskamp (2018), when estimating a population distribution, participants are apt to activate sample from the biggest category (i.e., the one with largest frequency) first and then proceed toward the one with the lowest frequency. However, since the frequency of the category activated first is normally underestimated, population estimates are higher for rare categories in the social circle and lower for frequent categories, compared to social-circle distributions (Galesic, Olsson, & Rieskamp, 2018).

[xxx] This pattern of judgemental errors is observed in many real-life scenarios. Take an example of estimates of national income distribution. Dawtry, Sutton, and Sibley (2015) report that wealthier Americans, compared to the less wealthy, estimate higher (mean) levels of wealth in their social circles and the U.S. population. Moreover, Proto and Sgroi (2017) show that in several distinct domains such as political stance, mobile phone purchases, hights and weights, etc., people at extremes are inclined to believe themselves to be closer to the middle of the distribution than is the case.

homophily while being uncertain about their own characteristics may be deluded into believing that they are in the categories that the modes of the population distributions occupy. As a result, if contexts are thought to reflect the population's true taste distribution, relative ranks of available options will be regarded as a credible signal of popularity that indirectly informs false-consensus decision makers of options they are supposed to choose. Given that many human characteristics and traits, such as IQ scores, heights, personalities, and performances, which strongly influence choices, are commonly presumed to be approximately normally distributed (O'Boyle Jr. & Aguinis, 2012; Sartori, 2006; Uysal & Pohlmeier, 2011), the rank-order decision rule, combined with the above speculation, predicts that false-consensus decision makers will perceive the middle options as optimal and thus behave as if they irrationally pursue the compromise options.

Finally, I return to the original principle of Wernerfelt's (1995) theory. Broadly speaking, the idea of matching embodies a special form of social influence, whereby contexts act as social norms that reveals information about the population's tastes for decision makers to infer their preferences. Put simply, apart from other types of social influence such as social conformity, compliance, identification, and peer pressures, the rank-order decision rule implies that affecting decision makers' inferred preferences via context is the main channel through which social norms shape choice.

Relating to the examples of increased portion sizes presented at the beginning of the chapter, the notion that choice context communicates information about social norms is reminiscent of the most widespread account of the portion size effect, although its emphasis is typically placed on the role of other kinds of social influence. A large body of research in the field of food choice (e.g., Herman, Polivy, Pliner, & Vartanian, 2015; Rolls, Morris, & Roe, 2002; Wansink, Painter, & North, 2005) has long argued that portion sizes in the market may subtly set a consumption norm and guide the amount that is perceived to be appropriate to eat. More importantly, many experimental studies demonstrate that portion sizes can vary substantially while still being considered equally normal and sensible for people to follow[xxxi] (Robinson, te Raa, & Hardman, 2015; Robinson, Oldham, Cuckson, Brunstrom, Rogers, & Hardman, 2016). For instances, one recent study conducted by Robinson, Henderson, Keenan, and Kersbergen (2019) shows that exposure to a smaller portion size of snack food leads participants to regard, on average, the smaller portion size as “normal” and consume less on later food intake. These studies consistently highlight the perceived link between market context and norms of appropriate eating, and the effect of this link on consumption choice. Accordingly, observed choice change, or choice reversal, may be due merely to a change of perceived norms resulting from context change. This gives rooms for firms to manipulate choice context (such as portion sizes on sale) to raise profits.

As noted above, the mainstream literature concentrates on discussing how market context affects choice through setting standards for what constitutes appropriate consumption. Far less research, unfortunately, delves into contextual influences on preferences inferred from the market[xxxii]. Yet, in fact, Wernerfelt’s (1995) theory can more or less account for the portion size effects on the basis of social norms. As implied by his theory, even though context keeps changing, as long as perceived relative standings remain unchanged, options with the same relative ranks will still be chosen. Of course, this explanation rests on the assumption that people trust social norm indicators (i.e., market contexts) more than their own feelings about products and thereby rationally calibrate absolute judgements of their own preferences accordingly in order to maintain the same estimate of relative positions. Furthermore, Wernerfelt’s (1995) theory provides an alternative, illuminating perspective on why people’s feelings of an option are found

---

[xxxi] The reasons behind the tendency to converge with social norms are generally thought to be that conforming to other people’s behaviour is mentally rewarding (Higgs & Thomas, 2016; Klucharev, Hytönen, Rijpkema, Smidts, & Fernández, 2009), social norms are believed to be reasonable and correct (Deutsch & Gerard, 1955; Higgs, 2015), and imitating others is a good way to smooth social interaction (Chartrand & Bargh, 1999).

[xxxii] A few papers (e.g., Bem, 1972; Higgs, 2015) mention that conforming to social norms is likely to trigger a change in self-perception and attitudes. But this psychological mechanism is obviously different from the implication of Wernerfelt’s (1995) theory.

to shift with its relative rank[xxxiii] as well as why participants tended to purchase a product with a larger portion following a person who first buys a large one, while the choice is significantly smaller in size when that person is overweight than being slim (McFerran, Dahl, Fitzsimons, & Morale, 2010).

In summary, this subsection introduced the rank-order decision rule proposed by Wernerfelt (1995) and expounded three possible psychological factors underlying the compromise effect. To recap, Wernerfelt (1995) explains the compromise effect by arguing that decision makers are unable to correctly employ the rank-order decision rule under the first-appearing context because they know that the first context does not capture the general population's tastes. The compromise effect mainly arises from suboptimal choice made in the first context. Based on recent experimental evidence on social sampling, the present thesis surmises that compromise behaviour is possibly precipitated by inaccurate estimate of population distributions, naturally arising from in-group homophily and/ or a propensity to remember and activate extreme samples. Despite these differences, all these accounts indicate that seemingly irrational behaviours can be a result of rational inferences about one's innate, context-independent underlying preferences.

1.3.5 Experience in decision-making

Most existing models of the compromise effect assumes uncertainty in underlying preferences and/or lack of knowledge about options' features, but often do not make predictions about whether the compromise effect reduces with experience. Yet, in everyday life, market involvement is not one-off – consumers normally face the same purchase tasks repeatedly. Investigating whether experience per se is able to help people learn their true underlying preferences to defeat contextual influences is thus of great practical importance.

I consider the empirical findings on this issue first. Much existing experimental evidence appears to confirm that behavioural anomalies fade out as experience of decision

---

[xxxiii] Sharpe, Staelin, and Huber (2008) find that a beverage is judged as smaller and more appropriate when its relative rank is second-largest as opposed to the largest.

making accumulates. Many experimental studies (e.g., Brouwer, 2012; Hoeffler & Ariely, 1999) find a positive correlation between choice stability and experience in the same domain. In similar vein, Coupey, Irwin, and Payne (1998), Erev, Ert, Plonsky, Cohen, and Cohen (2017) and Kramer (2007) observe that experienced choosers on average behave more consistently across problem frames, choice tasks and elicitation methods than do inexperienced choosers. In addition, the degree of the willingness to pay and willingness to accept gap is much smaller for experienced traders (List, 2003, 2004, 2011) and in repeated experimental markets such as repeated Vickrey (1961) auction experiments[xxxiv] (e.g. Cox & Grether, 1996; Shogren, Margolis, Koo, & List, 2001; Plott & Zeiler, 2005). Concerning context effects, extensive research (e.g., Brouwer, Dekker, Rolfe, & Windle, 2010; Duncan, 1972; Hoeffler, Ariely, West, & Duclos, 2013) finds that experience may equip decision makers with information and knowledge about the results of choosing certain options to overcome uncertainty and boost their confidence about choice, which are later proved to be helpful in purging the compromise behaviour[xxxv] (Chuang, Cheng, Chang, & Chiang, 2013; Chuang, Cheng, Chang, & Chiang, 2013). To sum up, such empirical observations seem to support the conjecture that experience gained in a repeated choice task can cause choosers to become more immune to external influences.

This conclusion is unsurprising from a standard economic point of view. The economic tradition tends to favour the notion that markets are institutions that help engagers learn and reveal their underlying preferences (Isoni, Brooks, Loomes, & Sugden, 2016). One relevant early representative theory is Plott's (1996) discovered preference hypothesis. This theory captures the idea that rational agents may utilise the market environment to acquire their underlying preferences by trial and error (Loomes, Starmer, & Sugden, 2003; Plott, 1996). Therefore, as experience increases, underlying preferences

---

[xxxiv] A Vickrey auction is a sealed-bid, second-price auction designed by Vickrey (1961), whereby bidders submit bids while being unaware of the bids of their opponents and the one who offers a highest bid wins but only needs to pay at the second-highest bidding price.

[xxxv] This may be because the extreme options are perceived as being riskier than others (Chernev, 2004), which scares off inexperienced or less self-confident people.

are expected to become more accessible[xxxvi] and more able to direct behaviour, promoting choice stability and consistency (List, 2003). To make it less vague, the hypothesis was later refined by Loomes, Starmer, and Sugden (2003), where the roles of repetitions, feedbacks and incentives in preference discovery are specified. By their supposition, repetitions enhance individuals' familiarity with choice tasks and options; feedback informs people of the consequences of certain decisions, and incentives encourage choosers to behave carefully during tasks. Notably, unlike Plott (1996), this refined version does not underline the involvement of the market mechanism since other choice environments that are characterised by these factors may exist. Hence, what this more developed hypothesis suggests is that as long as decision makers perform the same choice tasks repeatedly, with incentives and feedback regarding the results of previous trials, the outcomes will converge to an anomaly-free state (Loomes, Starmer, & Sugden, 2003).

Nonetheless, these are not the only interpretations of the disappearance of anomalies. This evidence can also be interpreted within an entirely different, and even opposite, perspective. Despite agreeing on the idea that some learning occur during repeated trials, the shaping hypothesis holds that people with *ex ante* blurry preferences have a natural tendency to adjust their inferred preferences in response to environmental cues[xxxvii], which need not contain relevant information (Loomes, Starmer, & Sugden, 2003, 2010). In other words, this hypothesis suggests that there are no such things as pre-existing, well-articulated underlying preferences waiting to be discovered. Rather, people generate an answer using heuristics to address a task (Loomes, Starmer, & Sugden, 2003). Along the same lines, Ariely, Loewenstein, and Prelec (2003) experimentally show that subjective valuations converge to an arbitrary, stable level, subject to contextual or other framing cues. This adjustment process effectively helps decision makers gradually behave more coherently and consistently given a fixed context (as if they really have "discovered" their underlying preferences). Accordingly, the conventional revealed preference approach in

---

[xxxvi] In other words, people are more aware of what they truly like and dislike.

[xxxvii] For example, consumers in repeated markets may refer to the observed market price to revise their subjective valuations of commodities.

which underlying preferences are estimated through observed data may be fallacious. In summary, these viewpoints imply that instead of inducing individuals to act on their underlying preferences, learning through (market) experiences moves behaviour towards context-specific patterns, contradicting the predictions of traditional economic theory.

To illustrate this set of ideas more concretely, take an example of the disparity between willingness to pay and willingness to accept. Based on the shaping hypothesis, the vanishing of the gap observed in repeated Vickrey auction markets (e.g., Shogren, Margolis, Koo, & List, 2001) may result from bids and asks being adjusted towards the market-clearing price presented in the previous periods[xxxviii], which is not necessarily affiliated to participants' true ideal values (Loomes, Starmer, & Sugden, 2003). Moreover, although the gap is found to be attenuated in a second-price auction, it can still be found in a strategically similar, second-to-last price auction[xxxix] (Knetsch, Tang, & Thaler, 2001). Thus, Loomes, Starmer, and Sugden (2003) conclude that decision-making experience *per se* may be unable to erode behavioural anomalies, and experimental evidence that appears to support those arguments may be an artefact of the experimental designs adopted.

Targeting the conflict between these two contrasting viewpoints, Amir and Levav (2008) develop a third, and more compatible, argument, in which learning strategy in

---

[xxxviii] In fact, the empirical results of this hypothetical market dynamics are mixed. Tufano (2010) provides evidence to support the shaping hypothesis, as an alternative to the discovered preference hypothesis, while Braga, Humphrey, and Starmer (2009) only find a tendency to lower bids posterior to unpleasant outcomes, rather than an inclination to calibrate bids in terms of market prices. This suggests that the market may shape consumer behaviour, but maybe not always in the way depicted by the shaping hypothesis.

[xxxix] Both types of auctions are modified versions of the Vickrey auction. According to Knetsch, Tang, and Thaler (2001), in the second-price auction, ten mug sellers (buyers) were told that the seller (buyer) willing to sell (buy) at the lowest (highest) price sells it at the price of the second lowest (highest) bid. By contrast, in the second-to-last price auction, mug sellers and buyers were instructed that eight mugs will be sold (bought) at the ninth lowest (highest) price (Knetsch, Tang, & Thaler, 2001).

repeated markets depends on the structure of the choice environment. Firstly, they find that a binary choice set with evident attribute trade-offs may impel decision makers to evaluate and learn their subjective valuations on each attribute, leading their final inferred preferences to be relatively stable across contexts. This learning propensity possibly results from the problem regarding trade-offs between conflicting attributes that people can only resolve by determining their subjective value for each attribute, as well as the absence of any effect of salient environmental cues (Amir & Levav, 2008). Contrarily, when learning in a trinary set, especially in an attraction and compromise paradigm, Amir and Levav (2008) find that people tend to learn context-specific heuristics[xl] rather than engaging in arduous value assessment[xli]. The authors then ascribe this tendency to the fact that an increased number of options creates a greater choice uncertainty that makes it even harder to learn subjective valuations. In contrast, more pronounced contextual cues may allow people to address the local choice tasks without relying on relatively vague underlying preferences (Amir & Levav, 2008). In other words, through interacting with choice contexts, decision makers learn how they prefer to deal with the local tasks, which keeps them from needing to figure out their willingness to trade off attributes. However, compared to trade-off learning, this type of learning strategy may produce stronger preference uncertainty, less sensitive valuations on an attribute, and more inconsistent choices as context changes (i.e., as if learning never takes place). In summary, Amir and Levav's (2008) findings demonstrate that people do not always attempt to learn subjective valuations of attributes or the underlying preference relations, which, as a guideline, are portable across contexts. In some decision-making structures, they may instead favour relying on contextual cues to develop a context-specific choice heuristic for problem solving.

[xl] For example, in a compromise condition, individuals learn to always select the compromise option, in lieu of seeking the one they may have an underlying affinity for.

[xli] In effect, in some circumstances, people may still engage in trade-off learning under large contextual influence. For example, experts may intentionally learn knowledge of attribute values and trade-offs and hence exhibit stronger resistance to environmental influence (Carlson & Bond, 2006).

Interestingly, decision making under risk reveals a different story about the impact of experience. Recent experimental results (e.g., Ert & Lejarraga, 2018; Hadar, Danziger, & Hertwig, 2018) show that context effects that have been reported in decisions-from-description studies cannot be observed in situations involving learning from experience[xlii]. Even worse, Spektor, Gluth, Fontanesi, and Rieskamp (2019) find reversals of the attraction and the compromise effects in decisions-from-experience scenarios, while a significant similarity effect is still observed. This indicates that the "standard" context effects may not be robust in decisions-from-experience environments. One possible reason is that, for context effects to emerge, the differences (or dominance relationships if any) between lotteries must be large enough for people to clearly recognise[xliii] (Hadar, Danziger, & Hertwig, 2018; Huber, Payne, & Puto, 2014; Simonson, 2014). This condition becomes difficult to meet when lotteries' payoff-relevant information is not explicitly presented, requiring people to infer it over trials (Frederick, Lee, & Baskin, 2014; Hadar, Danziger, & Hertwig, 2018). Furthermore, this experimental evidence seems to be better explained with reference to attentional and retrieval biases. As argued by Spektor, Gluth, Fontanesi, and Rieskamp (2019), the nature of decisions-from-experience tasks lead options with similar feedback value to inhibit each other in memory during evaluation, rendering both being perceived as less attractive. In contrast, options with distinct outcomes are more likely to stand out, receive more attention, and be

---

[xlii] Experimental paradigms in decision making research can often be classified into description-only and experience-only tasks, in which the former, to date, is far more popular in behavioural research than the latter is (Weiss-Cohen, Konstantinidis, Speekenbrink, & Harvey, 2016). Technically speaking, the distinction between these two paradigms concerns whether payoff-relevant information is exhaustively described (i.e., information about potential outcomes of choice alternatives and their associated probabilities are described for decision makes) or sampled (i.e., which outcomes will occur and what their probabilities may be have to be learnt through experience since there is no such conclusive information).

[xliii] In Amir and Levav's (2008) language, when differences among (three) options are clearly perceived, contextual cues become salient and play a dominant role in learning, reducing the inclination to learn absolute subjective values and willingness to trade off attributes.

perceived as more attractive, in line with a wealth of past studies on human attention (e.g., Krajbich, Armel, & Rangel, 2010; Madan, Ludvig, & Spetch, 2014; Theeuwes, 2010). The distortion of individuals' inferences about options' characteristics eventually drives decision makers to prefer options that are dissimilar to others (Spektor, Gluth, Fontanesi, & Rieskamp, 2019). This implies that cognitive biases may play a crucial role in learning one's own underlying preferences, causing inferred preferences and resulting choices to fail to follow true preferences. Such behaviour would be incompatible with the predictions of the discovered preference hypothesis and its refined version.

To conclude, subject to the structure of the choice set, experience gained in repeated decision-making tasks may change individuals' behaviour by helping people learn their context-independent underlying preferences, possibly with some biases, or learn context-specific choice strategies from salient contextual cues to cope with the decision problems they face. Amir and Levav's (2008) account predicts that the compromise effect will arise and even persist in repeated markets since the desire to simplify the tasks eclipses the desire to make an optimal decision as the main driver. This then raises the question of whether the compromise effect can be eliminated by consumers' experience, which is one of the research questions in the present thesis.

1.3.6 Relative Rank Theory (RRT)

Continuing the previous discussion, a model that is useful in understanding the role of experience in compromise behaviour is introduced in this subsection. The selected model is RRT, proposed by Brown and Walasek (2018). Note that RRT does not explicitly suggest a psychological mechanism for the effect of market experience on decision making. Rather, it holds that the expressed preference[xliv] that guides the final choice is

---

[xliv] In language of the RRT, the term "expressed preferences" denotes preference learnt through binary comparisons as well as preference that directs downstream behaviour. That is, Brown and Walasek (2018) do not distinguish inferred preferences from expressed preferences. To make the definition of the terms coherent in the thesis, the present subsection does not follow

constructed via ordinally comparing the consequences of actions in pair[xlv]. RRT is introduced here because its preference-learning framework will be used to build up the first proposed model of the thesis and some its major assumptions will be retained in the second model. An implementation of extended RRT model will be presented in next chapter, while the following explains RRT in more detail.

The key feature of RRT is that it accepts the assumption from standard economics of stable underlying preferences to account for individual differences, while injecting this idea with a behavioural claim that people cannot consciously access the strength of these underlying preferences[xlvi] and therefore must infer them based on past choices. Due to the

---

their nomenclature – inferred and expressed preferences are still treated differently when describing RRT.

[xlv] RRT suggests that it is the expressed (or inferred) preference, rather than the underlying preference, that directly determines final choices. However, this preference does not exist naturally – it must be formed by decision makers through a certain preference-learning process under a choice context. More concretely, in RRT, it is assumed that individuals can realise their preference over two alternatives only after physically experience them, and this preference relation should be ordinal and based on the underlying preference. Accordingly, RRT assumes that the expressed (or inferred) preference is constructed through a series of binary, ordinal comparisons between a pair of experienced choice alternatives. The whole formation process of the expressed (or inferred) preference and explanations of relevant assumptions will be presented in the rest of the section.

[xlvi] The term "cannot consciously access" means that the strength of the underlying preference can be processed to affect more complex cognitive processes such as decision making, but itself is unknown to decision makers. That is, despite having a pre-existing underlying preference that can be processed by the brain, individuals cannot retrieve, and thus speak out, their true, absolute valuations of options they evaluate or value representation of consumption utility. According to RRT, individuals can acquire no better than ordinal information about their underlying preferences over products after consuming or utilising them. Moreover, it is important to note that the term "consciousness" is rather ambiguous in the literature. For example, it is typically viewed as "wakefulness" in clinical practice, whereas

assumption of scale inaccessibility, preference judgements regarding available options are assumed to be formed by a series of binary, ordinal comparisons of paired options' consumption outcomes (Brown & Walasek, 2018). More specifically, during preference construction, the underlying preferences merely underpin the ordinal dominance relation for each compared pair, whereas they are incapable of revealing any cardinal information such as by how much one is preferred to another (Brown & Walasek, 2018). As a result, contrary to Plott's (1996) discovered preference hypothesis, in RRT the inferred preferences, that is the preferences that are inferred from outcomes of binary comparisons in a given context and are responsible for directly and exclusively guiding downstream behaviour, are defined in relative rank coordinates, not in absolute, real-world quantities. In other words, the inferred preferences can only be informative about the relative attractiveness of relative ranks of options. Brown and Walasek (2018) argue that this context-specific, rank-dependent mechanisms is the cause of seemingly irrational or biased behaviour[xlvii].

There are three key assumptions at the heart of the RRT. First, it assumes that there exists a context-independent, yet scale-inaccessible, underlying preference, which reflects

---

usually thought as "awareness" elsewhere (Cavanna, Shah, Eddy, Williams, & Rickards, 2011; Zeman, 2006). Roughly speaking, in terms of information processing, both "consciousness" and "awareness" can be defined as the content of one's subjective experience that may involve in human information processing and is composed of various psychological processes such as sensations, thoughts, emotions and memory (Zeman, 2001; Zeman, 2006). Here, however, when stating the assumption of restricted accessibility to the underlying preference, the thesis tends to use the term "consciousness" rather than "awareness" since the former is more often used together with human brain and reflection.

[xlvii] For example, people who have derived an inferred preference that suggests an option lying at the 40 percentile of the choice set is the mostly preferred one may tend to choose the option with the relative rank equalling to 0.4, even when context changes. This tendency will not be altered until they re-learn inferred preferences in the new context. Therefore, people may behave as if they have a preference towards a certain relative rank of products, instead of towards attribute values.

properties of a real-world object on an absolute psychological scale. As noted, the assumed existence of a stable underlying preference is conceptually consistent with the traditional economic account of preference, in that both perspectives view preference as an innate characteristic that is formed prior to any choice task and is stable within a person. Moreover, the assumption that the strength of an underlying preference is not consciously accessible accords with the long-lasting view of the “strangers to ourselves” perspective (Wilson, 2002) in social psychology. The theory suggests that decision makers have little or no direct access to their high-level cognitive processes (Keeling, 2018; Nisbett & Wilson, 1977) as evidence shows that their stated reasons for choices or ability to predict their future behaviours may be no better than a friend or stranger’s (Helzer & Dunning, 2012; Jackson, Connolly, Garrison, Leveille, & Connolly, 2015; Vazire & Mehl, 2008). Following the theory, RRT assumes that people have no conscious access to an internal subjective scale of an object and that therefore the scale cannot directly underpin the final choice.

The lack of privileged access to decision-relevant mental facts implies that people form their judgements possibly through a process of interpreting their own behaviour. Hence, RRT assumes that preferences are inferred from a comparison sample provided by individuals’ own past choices. Firstly, this assumption rests on self-perception theory (Bem, 1967, 1972), which argues that people come to realise their internal states, including preferences and attitudes, by inferring them from the history of their own choices, particularly when internal cues are too weak or ambiguous to be recognised. In other words, the theory suggests that people may adjust their internal values to align with *ex post* rationalisation of their choice (Bem, 1967). This resonates with the later-found phenomenon, namely choice-induced preference change, which shows that the post-choice subjective valuations of chosen items are usually higher than their pre-choice estimates (Izuma & Murayama, 2013; Sharot, Fleming, Yu, Koster, & Dolan, 2012; Sharot, Velasquez, & Dolan, 2010; Voigt, Murawski, & Bode, 2017). This even happens when the “choice” given to be *ex post* rated or justified is not decision makers’ intended one (Hall, Johansson, Tärning, Sikström, & Deutgen, 2010; Johansson, Hall, Sikström, & Olsson, 2005). In a nutshell, due to constrained conscious access to the underlying preference, decision makers are assumed to infer their preferences through their past choice.

Further, RRT assumes that the process of constructing inferred preference involves a series of ordinal, binary comparisons of choice samples in memory. Sample comparisons are assumed to be ordinal and binary since the underlying preference is assumed to underpin only ordinal choices between pairs of options that had been *ex ante* experienced. This assumption is inherited from the theory of Decision by Sampling (Stewart, Chater, & Brown, 2006), which asserts that mental judgements are formed based on pairwise, ordinal comparisons between the target option and other stimuli retrieved from memory. Also, the view that judgements are formed by a comparison of salient samples is compatible with a wealth of work in economics (Bordalo, Gennaioli, & Shleifer, 2013, 2020; Jehiel & Steiner, 2020), psychology (Denrell, 2005; Fiedler, 2000; Hertwig & Pleskac, 2010; Noguchi & Stewart, 2014; Ronayne & Brown, 2017), and neuroscience (Bornstein & Norman, 2017; Hu, Domenech, & Pessiglione, 2020; Shadlen & Shohamy, 2016). Notice that since the inferred preference is derived from a process of ordinal comparisons under a particular context, it must be defined in relative rank coordinates. This property may in turn account for the finding that people's subjective judgements of physical stimuli are often relative-rank based and context-sensitive (Aldrovandi, Wood, Maltby, & Brown, 2015; Boyce, Brown, & Moore, 2010; Niedrich, Weathers, Hill, & Bell, 2009; Riskey, Parducci, & Beauchamp, 1979). In summary, due to the limited cognitive accessibility of underlying preferences, decision makers are assumed to infer their preferences through ordinal, pairwise comparisons of experienced options. This rank-based inferred preference is the one that directly guides final choice.

The distinction between underlying and inferred (or expressed) preferences allows RRT to complement the literature by explaining "biased" decision making from the angle of context-specific preference learning while retaining the standard assumption of stable underlying preferences in economics (as is needed to capture individual differences). This idea echoes Isoni, Brooks, Loomes, and Sugden's (2016) proposition that attributing context-dependent behaviours to completely malleable underlying preference may be too extreme. In the light of this, the present thesis will adopt RRT's framework, particularly the process of pairwise comparisons, to build up the first model, and will retain some key assumptions in the second model.

1.4 Gaps, objectives, and thesis outline

Motivated by the observation of the upward shift in portion size and food intake, the present thesis poses three key questions regarding contextual influences on choice. The first question is about the existence of an innate, context-independent underlying preference that guides choice behaviour. If there is one, why and how do people's choices often appear to be susceptible to influence by market contexts? Second, it is asked under what conditions the (one-dimensional) compromise effect will or will not persist in market equilibrium when consumers have extensive market experience. Further, is there any other factor that can yield the compromise effect, so that experience *per se* is unable to completely remove contextual influences? Finally, the thesis asks how a profit-maximising firm reacts to consumer bounded rationality and its impact on consumer welfare.

There is no existing overarching model that addresses these questions simultaneously. Although models listed above tend to, explicitly or implicitly, assume stable underlying preferences, many of them do not discuss the role of experience in the compromise effect. It is speculated here that most models in the first three groups predict disappearance of the compromise effect in a repeated-purchase scenario since they generally ascribe the effect to uncertainty about preferences, attribute information, or choice outcomes. Yet, this makes those models unable to explain the trend in increased portion size and food intake. Moreover, most of these models do not explore factors that limit contextual influences. As to firms' responses, only a few models touch on this topic, with some demonstrating that firms may have an incentive to introduce a high-end option that is unprofitable on its own but that nevertheless prompts demand for other goods (e.g., Kamenica, 2008) while other models (e.g., Piermont, 2017) do not support this proposition.

In addition to the failure to answer the proposed questions, there are other common limitations of existing models of the compromise effect. Firstly, the involvement of attribute trade-offs in decision making is typically assumed as a default and is sometimes viewed as a key driver of the effect (e.g., Dhar, 1996; Simonson, 1989). The possibility that the compromise effect may exist in a one-dimensional case is often neglected in the literature. Secondly, many models work only when choice context is shifted from a binary to a trinary set (e.g., Guo, 2016; Kamenica, 2008; Wernerfelt, 1995). These models cannot easily be generalised to a scenario where the size of choice set is the same before and after contextual change. Thirdly, most conventional models do not capture external factors such as social influences, much less the effect of social norms on preferences inferred from the market offerings.

1.4.1 Summary of the proposed models

To fill this void, the present thesis goes beyond current theoretical models of the compromise effect by offering two alternative models based on earlier work and evidence, namely RRT (Brown & Walasek, 2018), the rank-order decision rule (Wernerfelt, 1995) and the false consensus effect (Mullen et al., 1985). Both of the proposed models inherit RRT's assumption of the distinction between underlying and inferred preferences to capture individual difference and common behavioural tendencies such as the compromise effect. Moreover, the idea that consumers learn their inferred preferences through physically experiencing and/ or observing their own choices is also retained from RRT. More specifically, the first model assumes that consumers perform a series of pairwise, ordinal comparisons to infer their preference before making final decisions. Here, the preference-learning stage and the decision-making stage are separate and one-way, unless people want to relearn their preferences. Thus, consumers in the first model will always choose according to the same inferred preference throughout all purchasing periods. In contrast, the second model relaxes the assumption of this one-way path by assuming that consumers may choose sometimes based on past choice records, i.e., their inferred preferences, and sometimes on the "norm-based" preference (defined below) or randomly. This potentially leads choice and the inferred preferences in the following rounds to vary across choice rounds.

In both models, norm-based preferences embody the influence of perceived social norms on people's inferred preferences. The specific assumption of the norm-based preference results from the rank-order decision rule, which suggests that people tend to choose the product whose rank in the set matches their perceived relative standing in the population's taste distribution. Both models assume that the inclination to believe that one is at the middle of the taste distribution biases the norm-based preference towards favouring the middle options.

1.4.2 Contribution

With the proposed models and their predictions, the present thesis contributes to the literature by making the following arguments. Firstly, attribute trade-off is not necessary for the compromise effect to occur as both models predict that difference in values along one dimension is sufficient to trigger it. Secondly, the compromise effect may still exist even in a market equilibrium in which consumers can repeatedly purchase and learn their preferences through experience. This further implies that the seemingly irrational behaviour represented by the compromise effect can be a result of a rational, albeit biased, inference about one's own underlying preferences. This also highlights the importance of the ability to accurately estimate one's own position in the population as well as to identify deceptive information conveyed by market context. Thirdly, it is posited that the existence of stable underlying preferences is a crucial factor in preventing market context from being unlimited. These internal cues may serve to help consumers detect relative (un)desirability of options. Fourthly, context manipulations by the firm will happen when it is profitable to do so. This will potentially shrink consumer welfare as distorted menus mislead consumers more than they educate them.

1.4.3 Thesis Outline

The remainder of the thesis is organised as follows. Chapter 2 and 3 will respectively present models of the compromise effect with preference learning to address the research questions outlined above. Both chapters will consist of a formal description of the model and the market along with several computer simulations. Chapter 4 will present an experimental study that uses a non-standard approach to test key assumptions adopted in the proposed models such as the false consensus effect and the rank-order decision rule. The thesis will close by discussing theoretical and practical implications, along with limitations and recommendation for future research.

Chapter 2: Model 1

Focusing on the compromise effect, the chapter begins to address the three questions, identified in Chapter 1. To recap: The first question concerns the existence of context-independent underlying preferences[xlviii]. Irrespective of the empirical data, it is of theoretical importance to understand why and how people with context-independent underlying preferences may still exhibit a compromise behaviour. The second question then asks about the conditions under which the compromise effect will or will not persist in market equilibrium when consumers have an opportunity to learn their underlying preferences by exploring all market offerings. To answer this question, it is important to investigate whether there are other factors that lead to the compromise effect. These two questions arise because, in reality, consumers appear not to choose the middle option under all market contexts, while exhibiting the compromise behaviour under some contexts despite enormous experience. Therefore, it is essential to identify in advance what other factors contribute to the compromise effect if experiencing all market offerings is possible. This chapter examines the informative role of market context and social norms in compromise behaviour.

The last question relates to a producer's strategic reactions. In markets where changing menus is profitable (e.g., in the soft drinks market, a firm can make more profits by increasing the portion size), what is the firm's optimal menu-setting strategy in response to consumers' compromise tendency? Moreover, the question of how market equilibrium and consumers' welfares are affected by distorted product lines is also of great concern.

[xlviii] As defined in Chapter 1, the "underlying preference" refers to the preference that is inborn (i.e., biologically determined) and therefore is context-independent and pre-existing prior to a choice task. By contrast, the term "inferred preference" denotes a preference that individuals learn and form assumably from experience consuming choice alternatives during the process of choosing.

The lack of an overarching model that tackles all these questions at once points to the need to develop an alternative, more comprehensive model of context effects. The present chapter therefore aims to begin development of a plausible psychological mechanism that underlies consumer behaviour, including learning, to explain why market experience may result in compromise behaviour and at the same time why it does not eventually, given firms' incentives, push attribute values of market offerings to infinity.

The chapter begins by explaining the rationale behind the model, followed by a formal description of consumer and producer behaviour in the market. Then, several computer simulations are presented. The simulations first explore a single-consumer case and then move to a scenario with multiple, heterogenous agents. The chapter ends with a summary of simulation results as well as their implications and limitations of the model.

2.1 The model

2.1.1 Overview

The first model is built on three theories: RRT (Brown & Walasek, 2018), the rank-order decision rule and the matching strategy (Wernerfelt, 1995), and the false consensus effect (Marks & Miller, 1987). The framework of the model is based on RRT. Many assumptions and rules made in the RRT are retained, including the distinction between underlying preferences (preferences$_U$) and inferred preferences (preferences$_I$), characteristics of these two preferences, the way people infer their preferences$_U$, and the criterion for choice deferral. Inspired by Wernerfelt (1995), the model introduces a new type of preference, namely preference$_N$ (norm-based preferences), to RRT to capture the idea that consumers usually treat menu and social norms as informative and tend to use them to assist decision making. This extension allows the model to answer the thesis' research questions.

In line with RRT, the model assumes that, while individuals are endowed with an underlying preference, which is defined in terms of context-independent quantities, its absolute magnitudes are not consciously accessible [xlix]. Therefore, preferences$_U$ can underpin only ordinal choices between pairs of alternatives that have been experienced. In other words, over a pair of options A and B, people can state that they prefer A to B or B to A after experiencing both of them, but are unable to indicate how much they prefer one to another. This restriction requires people to use some decision rules to make a choice in a set that contains more than two options.

According to RRT, when faced with a choice set people start learning their (rank-based) preference for each option through a series of binary ordinal comparisons, and the choice made in each binary comparison is purely guided by their preferences$_U$. The model proposed here differs from RRT in that the outcome of each pairwise comparison is also influenced by environmental factors such as menus. This idea comes from Wernerfelt's (1995) theory of market inference. In his theory, consumers often assume that firms make rational product-design choices given the preference distribution of the population, and that therefore, when consumers are uncertain about their true, absolute valuations of products, they may rationally choose the product with a rank that corresponds to their perceived rank in the taste distribution. For instance, if a consumer believes he/she is at the 20th percentile of the taste distribution in the population, he/she would choose the 20th percentile option from the ordered set.

[xlix] As noted in Section 1.3.6, the term "cannot consciously access" means that the absolute values of the preference$_U$ over an attribute can be processed by the brain to affect judgements and decision making, without being aware by decision makers. Therefore, individuals cannot retrieve, and thus precisely report, their innate, absolute valuations of choice alternatives they evaluate. In addition, the term "consciousness" and "awareness" have similar definitions in terms of information processing and thus often used interchanably in the field of behavioural science. More specifically, both refer to the content of individuals' subjective experience that may involve in human information processing (Zeman, 2001; Zeman, 2006). That being said, when it comes to describing the accessibility of the underlying preference, the thesis tends to adopt the term "consciousness" since it is more commonly used in depicting things related to human brain and reflection.

The results of market inference coincide with the predictions from standard economics only when consumers know their true relative standing of the population's taste distribution and also menus are designed to reflect that distribution. Indeed, a host of studies such as those of Burson (2007), Gershoff and Burson (2011), and Hamilton, Ratner, and Thompson (2010) find that many consumers do not make optimal purchase decisions because they typically have incorrect beliefs about their relative standing. For example, the false consensus effect reveals that many people overestimate the extent to which their own attitudes are shared by the population at large (Hoch, 1988; Marks & Miller, 1987). This highlights the important role of precision of estimation about one's relative position in the population in determining the accuracy of market inferences. Given the idea of market inference, consumers who inaccurately perceive themselves as an average person will be more likely to choose the middle option. Consequently, the presence of false-consensus consumers contributes to the observation of compromise behaviour in the market.

More concretely, the learning process in the model operates as follows. A figure that visually presents the learning process is displayed in the next subsection (Figure 1). Firstly, consumers are assumed to hold no prior knowledge of products, but are endowed with stable underlying preferences, which they cannot consciously access. After seeing all market alternatives, consumers start inferring their preferences through a series of binary, ordinal comparisons. For each pair of market options, one option is chosen over another only when its preference$_U$ value exceeds that of another by a certain amount[1] (the threshold criterion). Otherwise, no choice is made for that pair. Consumers are therefore not forced to make a choice during pairwise comparison. Conversely, if the difference in preferences$_U$ passes the threshold criterion, the actual choice probability of a paired alternative will be governed by both internal signals (preferences$_U$) and market inference

[1] The relative comparison decision rule is also inherited from RRT. The use of it captures the findings that undesirability of options (Dhar & Sherman, 1996; Gerasimou, 2017; Mochon, 2013) and cognitive and emotional difficulty of selecting (Anderson, 2003) are important drivers of choice deferral.

(the rank-order decision rule). Since "false-consensus" individuals may overestimate the extent to which others share their attitudes and opinions, at this stage the option whose rank in an ordered choice set is closer to the middle will be more likely to be selected from the pair.

After experiencing all market options and comparing them in pairs, consumers are assumed to infer their preferences (their preferences$_I$) by computing how frequently a product has been chosen relative to other market alternatives. Note that in both RRT and the proposed model, preference$_I$ are defined not in terms of absolute, real-world quantities, but in relative rank coordinates. That is, consumers can only know that they prefer a product at a given percentile of the distribution of choice set offered in the market. Finally, these learned preferences$_I$ drive purchase decisions in all following consumption periods (in which consumers are assumed to always choose the product with highest preference$_I$ value). Note that since the learning stage and the decision-making stage are separate, the preferences$_I$ remains unchanged across purchasing periods in this model unless consumers relearn their preferences.

Overall, the model represents and explores the possibility that the compromise tendency arises from biased preference learning in the market due to the false-consensus effect, rather than from unstable preferences$_U$, uncertainty about products, or attribute trade-offs. Moreover, it shows that if the value of the preference$_U$ of the middle option is too low to pass the threshold criterion when it is compared with other options, that option will not be chosen during learning, reducing compromise behaviour in the decision-making stage. More specifically, there are two conditions that must be met for the middle option to be chosen more frequently than others during learning. Firstly, its corresponding preference$_U$'s magnitude should sufficiently differ from that of other options. Otherwise, pairs that contain the middle option cannot pass the threshold criterion, rendering the middle option unlikely to be chosen during learning[li]. Secondly, if the middle option is

[li] More specifically, if the middle option is the second or third best option and its associated preference$_U$ value is not high (low) enough to make the pair that consists of it and an inferior (superior) alternative pass the threshold criterion, it will be less likely to be chosen,

not the optimal one, it should not be less preferred to dominant options by too much. Otherwise, it may not be selected from the pair as often as the dominant one is, even though it gains the largest benefits from the false-consensus effect. These two conditions together imply that the effect of false consensus on preference learning is restricted in the proposed model, and that preferences$_U$ serve to inhibit the possible extent of a firm's manipulation of market context.

2.1.2 Model description

This section introduces the psychological model of context effects in formal terms. There are four types of preferences assumed in the model, each of which has its own properties and functions. These are preferences$_U$, namely the predetermined, stable underlying preferences for consumers, preferences$_N$, namely the preferences obtained from market inferences, combined preferences (preferences$_C$), namely the combined preferences of preferences$_N$ and preferences$_U$, and preferences$_I$, namely the preferences formed by consumers' inference about their own preferences of available products. As argued before, the distinction among these four preferences is crucial in investigating the compromise effect with learning and stable preferences$_U$. A diagram that demonstrates the model's approach to the consumer decision-making process is displayed below (see Figure 1).

---

compared to the inferior option, during the whole pairwise learning process. This is because options with lower preference$_U$ values can still gain some, albeit very small, choice probability when being paired with the optimal option if the difference in the preference$_U$ value is high enough to pass the threshold criterion.

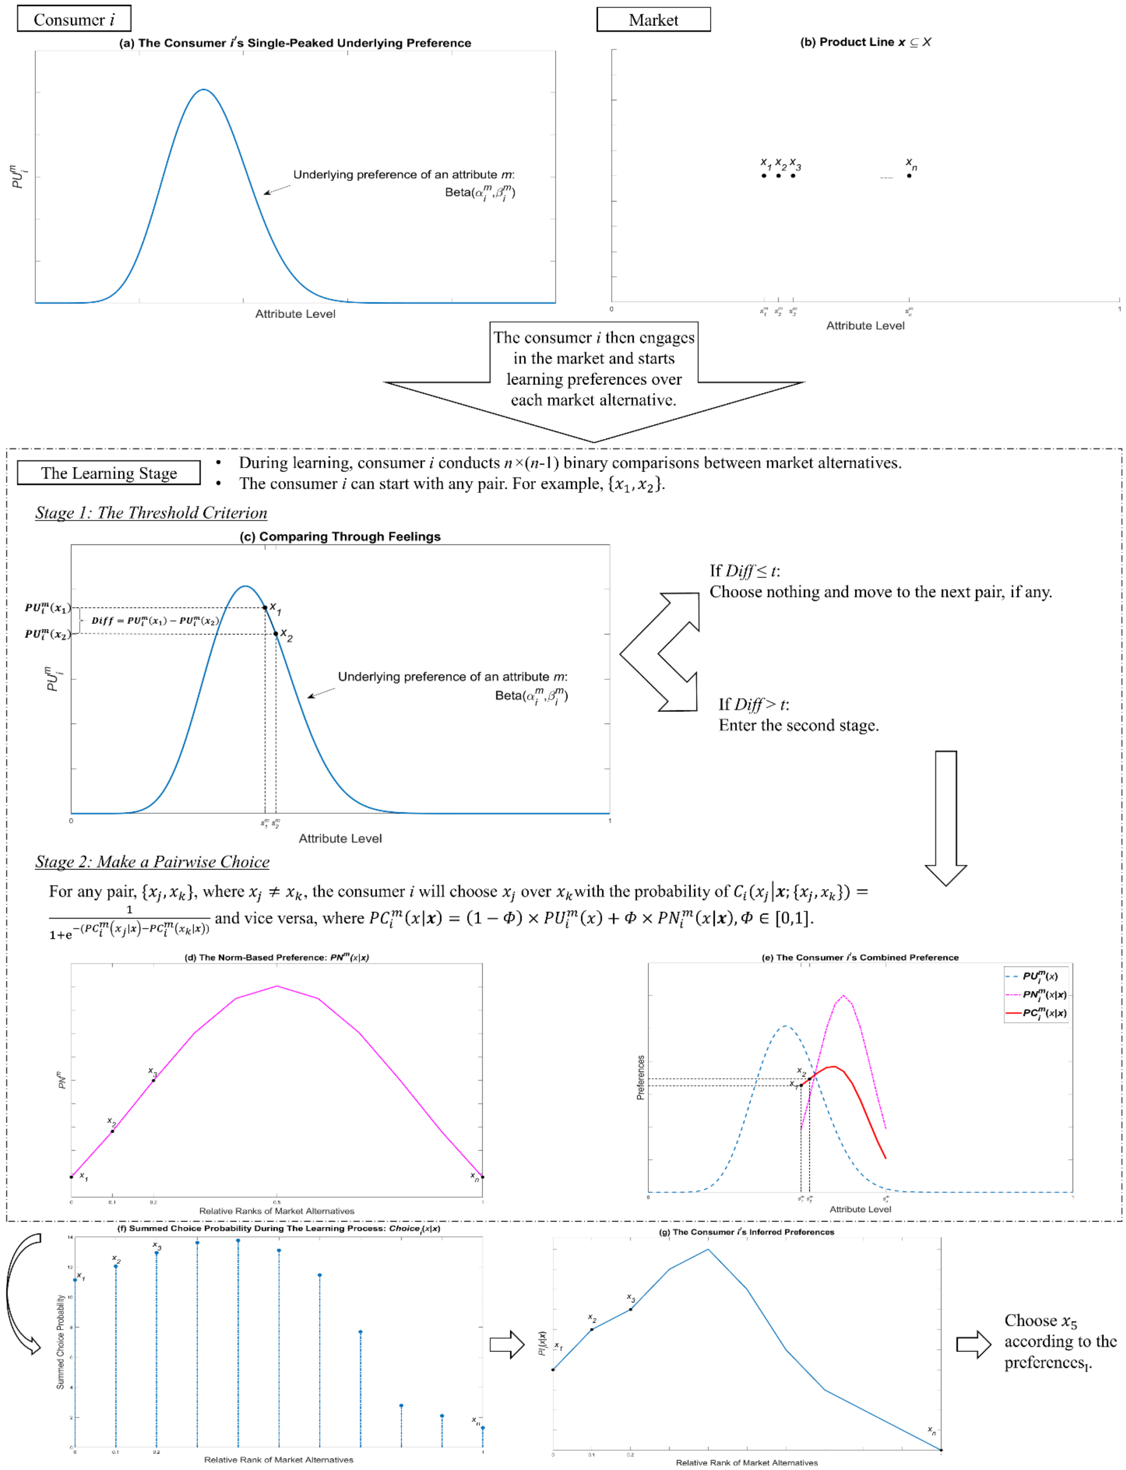

Figure 1. A graphical illustration of the model. Parameter values used in the figure are: n = 11, $\boldsymbol{s^m}$ = {0.36, 0.38, 0.40, …, 0.56}, α_i^m = 12, β_i^m = 24, Φ = 0.5, and t = 0.3.

2.1.2.1 The market

This subsection introduces the notations and assumptions about market context, which corresponds to the top-right corner of Figure 1. Consider a monopoly market where a profit-maximizing firm sells a range of products varying in attribute levels (e.g., quantity, quality, price) to a continuum of consumers. Let each $x \in X$ be a market alternative, and X be a compact set of technologically feasible market alternatives. Each alternative x is characterized by a vector of physical attributes, $\boldsymbol{s_x} \in S$.

Suppose the monopolist designs a product line consisting of *n* market alternatives $\boldsymbol{x}$ = {x_1, x_2, …, x_n} ⊆ *X*, in which all products are presented in ascending order with respect to the magnitude of a common attribute *m*, indexed by *j*, where *j* = 1, 2, …, *n*. A set of products' levels of the attribute *m* is hence written as $\boldsymbol{s_x^m} = \{s_1^m, s_2^m, \dots, s_n^m\} \subset \mathbb{R}_+$, with s_j^m representing the attribute level of x_j on *m*.

2.1.2.2 Preferences$_{\mathrm{U}}$

Now consider the demand side of the market where consumers have unit demand. The properties of preferences$_{\mathrm{U}}$ are described in this subsection. Suppose that for any attribute type *m*, the consumer *i* is equipped with a stable, single-peaked underlying preferences which satisfies some standard axioms of preferences such as completeness and transitivity. Following RRT, preferences$_{\mathrm{U}}$ over the full set of feasible quantities of an attribute *m*, $\boldsymbol{s^m}$, are represented by the probability density function of a beta distribution, $\mathrm{Beta}(\alpha_i^m, \beta_i^m)$, where $\alpha_i^m, \beta_i^m > 1$ to guarantee a unimodal distribution and quasi-concavity. The parameters α and β jointly characterize the distribution, thereby determining the precision of preferences$_{\mathrm{U}}$ and the absolute preference level over each attribute value. Based on its properties, the function of preferences$_{\mathrm{U}}$ is denoted as $PU_i^m: s^m \rightarrow \mathbb{R}_+$. An example of consumer i's preferences$_{\mathrm{U}}$ is shown on the top left of Figure 1.

Importantly, the exact value of any PU_i is assumed to be unknown by choosers. This assumption of inaccessibility of the strengths of preferences$_{\mathrm{U}}$ implies that despite having physically exploring all products in the market, *ex-post* knowledge on $PU_i(x)$[lii] is limited. As a result, preferences$_U$ can underpin only ordinal choices between alternatives with respect to one attribute type.

Put differently, let $\succsim_m$ be the perceived binary preference relation over the common attribute *m*, such that $\forall x, y \in X, x \succsim_m y \Leftrightarrow PU_i^m(x) \geq PU_i^m(y)$. For any choice pair $\{x_j, x_k\} \in \boldsymbol{x}$, the consumer *i* is able to state $x_j \succsim_m x_k$ and/or $x_k \succsim_m x_j$ after experiencing both choice options, while still being unaware of how much one is preferred to another or the exact magnitude of $PU_i^m(x_j)$ and $PU_i^m(x_k)$. Hence, if the taste difference is too small to be recognised, the consumer *i* may be incapable of establishing a preference ordering $\succsim_m$ based on internal feelings, violating the completeness axiom of revealed preferences.

2.1.2.3 The learning process: An overview

Resting on the above assumptions, the following model of consumer learning and decision-making process is assumed. Recall that consumers are *ex ante* inexperienced such that they have no prior knowledge of available products, $\boldsymbol{x}$, the technologically feasible set that subsumes those products, *X*, or their absolute preferences over the common attribute, *m*, $\boldsymbol{PU}^m$. They therefore need to learn their preferences in the market in order to maximise their expected benefits of consumption. The term "preference learning" in this model refers to the process whereby consumers conduct a series of binary ordinal comparisons to form an *ex-post* rank-based belief about their true tastes of products. This learning outcome then serves as a guide for final decision making.

2.1.2.4 The learning process: The first stage

The proposed learning process consists of two sequential stages, each of which is presented in Figure 1. This subsection formally introduces the first stage of the learning process.

[lii] For simplicity, $PU_i(x)$ is used to refer to the consumer *i*'s preferences$_U$ of each market alternative *x*, although PU_i is not literally a function of *x*. More specifically, $\forall j \in \{1,2,\ldots,n\}, PU_i^m(x_j) \stackrel{\text{def}}{=} PU_i^m(s_j^m)$.

In the first stage, consumer *i* compares pair alternatives along a common attribute, *m*, solely based on preferences$_U$, and then decides whether to make a choice from the pair. If the decision is "yes", they will proceed to the next stage at which they decide which pair alternative to choose. For *n* market alternatives, the consumer *i* will follow the same process for each pair till all $P_2^n = n(n-1)$ pairs are compared. Notice that the order of the elements in the pair matters, such that $\{x_j, x_k\} \neq \{x_k, x_j\}\ \forall\, j \neq k$.

The first stage of learning includes scenarios in which consumers may experience difficulty in choosing between pair alternatives due to a lack of a preference-dominant option. A threshold criterion that implements this consideration is introduced as follows.

For any pair, $\{x_j, x_k\}$, where $x_j \neq x_k$, the consumer *i* will proceed to the second stage of the pairwise comparison if and only if $PU_i^m(x_j) - PU_i^m(x_k) > t$ or $PU_i^m(x_k) - PU_i^m(x_j) > t$; otherwise, they will not make a choice from this pair and will move to the next pairwise evaluation, if any.

The threshold parameter, $t \in \mathbb{R}^+$, is viewed as a fixed dispositional trait. This threshold criterion implies that for a "reasonable" market context[liii], the more imprecise the preferences$_U$ are (i.e., the larger the variance of the preferences$_U$ distribution), the less likely a consumer is to make a choice, yielding lower consumer demand overall.

2.1.2.5 The learning process: The second stage

Conversely, if the taste difference is large enough to pass the criterion, the consumer *i* will need to decide which alternative is to be chosen at the second stage. This subsection provides a detailed description of the second stage of the learning process, including how preferences$_U$ and preferences$_N$ jointly determine pairwise choice.

The model extends RRT's approach according to which the choice made from each pair is no longer solely based on individuals' internal tastes. The model assumes that the uncertainty in subjective valuations of products may predispose people to rationally

---

[liii] The word "reasonable" means the set of attribute values of all market offerings does not lie within the area where the preference$_U$ is flat. In other words, a reasonable context refers to one that contains at least one option that is more desirable than others given the preference$_U$.

leverage external information to aid in decision making. More concretely, as suggested by Wernerfelt (1995), consumers tend to buy a product with a rank (within the market distribution) that corresponds to their estimated position in the population's taste distribution, presupposing that consumers believe the firm makes a rational product-design choice to meet the needs of all consumers. Accommodating this idea, consumer choice at the second stage of learning relies on two disparate devices – the preference$_{\mathrm{U}}$ and the preference$_{\mathrm{N}}$ – that operate in parallel and compete to control actual choice.

Formally speaking, let $PN^m(x|\boldsymbol{x})$ be a one-dimensional function that converts the preference ordering acquired from the rank-order decision rule into a real value, given the information that the firm introduced the product line $\boldsymbol{x}$. The magnitude of PN^m is represented by the probability density function of a normal distribution, $\mathcal{N}(\mu, \sigma^2)$. The values of the normal distribution parameters are estimated from $\boldsymbol{s}_{\boldsymbol{x}}^{\boldsymbol{m}}$, where μ is the mean of $\boldsymbol{s}_{\boldsymbol{x}}^{\boldsymbol{m}}$ and σ^2 is an unbiased estimator of the variance. Given the nature of the normal distribution, the false consensus effect is captured since the bell shape of PN^m implies that middle options are most attractive, whereas extreme options are least attractive. Mathematically, for any middle options, denoted by x_{Mid}, and $\forall j, k \in \{1,2, \dots, n\} \setminus \{Mid\}$, $\mathrm{k} \geq j > Mid \vee \mathrm{Mid} > j \geq k \Longrightarrow PN^m(x_{Mid}|\boldsymbol{x}) > PN^m\left(x_j|\boldsymbol{x}\right) \geq PN^m(x_k|\boldsymbol{x})$.

The combined preferences (the preferences$_{\mathrm{C}}$), PC_i^m, that determine the consumer i's preference ranking over a pair in the learning stage are defined as a convex combination of the two preference distributions:

$$PC_i^m = \{PC_i^m(x|\boldsymbol{x}): PC_i^m(x|\boldsymbol{x}) = (1 - \Phi) \cdot PU_i^m(x) + \Phi \cdot PN_i^m(x|\boldsymbol{x}), \forall x \in \boldsymbol{x} \subseteq X, \Phi \in [0,1]\}.$$

The parameter Φ is construed as the tendency for the consumer *i* to employ the rank-order decision rule *vis-à-vis* preferences$_{\mathrm{U}}$, with Φ = 0.5 indicating a balance of preference. A high Φ models the case where pairwise choices are more norm-based rather than feeling-based, whereas a low Φ represents a situation where feelings greatly outweigh the influence of market inference. Importantly, Φ also reflects the degree to which consumers trust market information. If consumers are convinced that the product line is randomised or distorted so that it does not contain any information that the monopolist possesses, they may assign Φ = 0 and the outcome of learning will then resemble the one in RRT.

Note that the relative precision of preferences$_U$ and preferences$_N$ may affect their relative contribution to PC_i^m, regardless of the value of Φ. More specifically, even when attention is paid equally to these two preferences (i.e., $\Phi = 0.5$), their relative effects on choice will depend on their relative magnitudes. The rationale behind this is straightforward. Suppose that one preference is more imprecise than another and that therefore, for some choice pairs, it may show relatively smaller differences in consumption utility between alternatives than the latter does. That is, for some pairs, the less precise preference$_U$ may indicate that pair alternatives are somehow indifferent to each other in terms of pleasure of consumption, whereas the latter is able to clearly represent taste differences between the two alternatives. Thus, when judging those pairs, the relatively more precise preferences will have more impact on the resulting decision since it offers clearer information on preference orderings.

Based on these assumptions, at the second stage of pairwise comparison, the choice probability is computed using a softmax function with the preference$_C$ as an input. Mathematically, the probability of choosing an alternative, x_j, over its paired option, x_k, under the context $\boldsymbol{x}$ is $C_i(x_j|\boldsymbol{x};\{x_j,x_k\}) = \frac{1}{1+\mathrm{e}^{-(PC_i^m(x_j|x)-PC_i^m(x_k|x))}}$. By the same token, the choice probability of x_k is $C_i(x_k|\boldsymbol{x};\{x_k,x_j\}) = \frac{1}{1+\mathrm{e}^{-(PC_i^m(x_k|x)-PC_i^m(x_j|x))}}$, or equivalently, $1 - C_i(x_j|\boldsymbol{x};\{x_j,x_k\})$. This means that no matter which pair option is indeed preferred in terms of preferences$_U$, both can generate some choice share as long as the threshold criterion for this pair is passed.

<u>**2.1.2.6 Post-learning stage**</u>

After performing comparisons of all possible pairs, the consumer *i* will make an inference about his/ her preference of each market alternative (i.e., preferences$_I$) based on the relative choice frequency of each product obtained from pairwise comparisons. This subsection formally describes how preferences$_I$ are formed and how they underpin final choice. An illustrative example is displayed at the bottom of Figure 1.

The summation of the probability that a market alternative x is chosen during preference learning is calculated as $Choice_i(x|\boldsymbol{x}) = \sum_j C_i(x|\boldsymbol{x};\{x, x_j\}) + \sum_j C_i(x|\boldsymbol{x};\{x_j, x\})$, where $x_j \in \boldsymbol{x} \setminus \{x\}$. Given this summed choice probability, the consumer *i* will make an inference about her personal preferences$_U$ by computing how frequently an option is chosen relative to other options in the learning stage. This inference about the preference of available products is named as the inferred preference (the preference$_I$), PI_i, and is defined as:

$$PI_i = \{PI_i(x|\boldsymbol{x}): \mathbb{Z}^3_{\geq} \to \mathbb{R}_+ \left| PI_i(x|\boldsymbol{x}) = \frac{\text{nlower}_x + 0.5\text{nequal}_x}{\text{nlower}_x + \text{nhigher}_x + \text{nequal}_x}, \forall x \in \boldsymbol{x} \subseteq X \right.\},$$

where $nlower_x$ indicates the number of alternatives chosen less frequently than option *x*, $nhigher_x$ refers to the number of alternatives chosen more frequently than *x*, and $nequal_x$ denotes the number of alternatives chosen equally often to *x* (*x* itself is not included). Note that PI_i is defined not in terms of absolute quantities, but in relative rank coordinates. That is, consumers can only know that they prefer a product at a given percentile of the product line. They do not know their preferences in terms of the absolute quantities of an attribute *m*, $\boldsymbol{s^m}$.

The *ex-post* learned preferences$_I$ will inform the final purchasing decision in the decision-making stage, given that the income constraint is not binding[liv]. Consistent with the standard assumption of utility maximisation, the unit-demand consumer *i* will buy the product with the highest PI_i. If more than one option has highest PI_i, those options will be chosen with equal probability[lv]. The consumer *i* will choose an exogenous outside option, *x₀*, if the shape of PE_i curve is completely flat.

[liv] This presumption rules out any price effects, and thereby allows the analysis to throw light on contextual influences on purely preference-guided consumer behaviour.

[lv] The model assumes that options in the set of *h* most preferred options will be chosen with probability $= \frac{1}{h} + \varepsilon_{i,\boldsymbol{x}}$, where $\varepsilon_{i,\boldsymbol{x}}$ are independent and identically distributed shocks with $\mathrm{E}[\varepsilon_{i,\boldsymbol{x}}] = 0$. On average, the probability of choosing an option in the set will approximate to $\frac{1}{h}$.

Moreover, it is assumed that once the firm changes the menu, consumer i will re-construct her preferences$_{\mathrm{I}}$ through the same learning process. Since the model assumes complete learning, past contexts will not influence consumers' evaluations in a new context at all.

Finally, recall that consumer choices change in response to the product line $\boldsymbol{x}$, rather than price. To clarify, in the rest of the thesis the notation, $D(x_j; \boldsymbol{x})$, will be used to depict the relationship between context $\boldsymbol{x}$ and quantity demanded of good x_j, holding the number of market offerings n constant[lvi]. More specifically, $D(x_j; \boldsymbol{x})$ specifies how many units of x_j are demanded in the various contexts, *ceteris paribus*. One noteworthy point is that for the same j, s_j^m may differ across contexts. Therefore, market demand $D(x_j; \boldsymbol{x})$ does not describe quantity demanded for a certain attribute level in different contexts.

2.1.3 The firm's profit-maximisation problem and the market equilibrium

This section analyses the monopolist's best response to market demand. To explore the firm's optimal strategy, the following will discuss the profit-maximising problem under two distinct cases: (1) products' marginal costs and prices are independent of attribute values, e.g., cheese varying in sharpness, and (2) products' marginal costs and prices linearly increase[lvii] with their attribute levels, e.g., soft drinks varying in portion size.

Let p_j and c_j be the price and marginal cost of the product x_j, respectively. To increase the transparency of the proposed underlying mechanism, it is assumed that p_j

---

[lvi] In order to ensure that the relative positions of products and their labels are consistent across contexts, products with same ranks in different contexts receive the same contextual influences.

[lvii] The assumption of linearity is necessary for one-dimensional products. Otherwise, there may involve a price-quality trade-off, which does not fit the research object of the thesis.

and c_j are exogenously determined[lviii]. Moreover, the fixed costs *FC*, which provide exogenous variation in the market context, are assumed to remain unchanged for a certain range of output levels.

Denote $\pi(x_j; \boldsymbol{x})$ as the monopolist's profit from selling x_j under a particular context $\boldsymbol{x}$. The total profit gained under the market context $\boldsymbol{x}$ is defined as $\Pi \equiv \sum_{j=1}^{n} \pi(x_j; \boldsymbol{x}) = \sum_{j=1}^{n} [p_j \cdot q_j^D(\boldsymbol{x}) - c_j \cdot q_j^s(\boldsymbol{x})] - FC$, where $q_j^D(\boldsymbol{x})$ and $q_j^s(\boldsymbol{x})$ refer to quantities sold and quantities supplied of option x_j under the context $\boldsymbol{x}$, respectively. Suppose that the monopolist has reliable information about demand for available market offerings, i.e., suppose that $D(x_j; \boldsymbol{x})$ is known. To obtain the maximum possible profit, the firm first identifies its profit-maximising supply function, $S(x_j; \boldsymbol{x} | D(x_j; \boldsymbol{x}))$[lix], which specifically gives the profit-maximising output level of x_j under feasible contexts, given consumer demand $D(x_j; \boldsymbol{x})$, and then choose that product line $\boldsymbol{x}$ that will lead to maximum profits. That is, the firm will use information about $D(x_j; \boldsymbol{x})$ to determine optimal output (supply) levels for each context, and then choose the context that maximises total profits, given consumer demand and the optimal output levels..

Selection of profit-maximising context involves determining the number of options, *n* and attribute levels of all products $\boldsymbol{s_x^m}$. To simplify the analysis, it is assumed that the values of the attributes *m* of each product form an arithmetic progression, whereby products' attribute levels differ from those of their adjacent neighbours' by a common difference, $d \in \mathbb{R}_{++}$. In other words, $s_{j+1}^m - s_j^m = d \; \forall \; j = 1, \dots . n - 1$.

[lviii] The assumption of exogenous price is necessary to guarantee a perfectly price-inelastic demand. If prices were endogenous, the monopolist would charge prices out of the reasonable range, binding the budget constraints. Consequently, the discussion of the pricing strategy, price effects, etc. would complicate the analysis.

[lix] In other words, $S(x_j; \boldsymbol{x} | D(x_j; \boldsymbol{x}))$ gives the optimal quantity supplied of the product x_j in all possible contexts, given the market demand $D(x_j; \boldsymbol{x})$. Similarly, to market demand, it portrays a relationship between profit-maximising output levels of x_j and context $\boldsymbol{x}$.

2.1.3.1 Case 1: Constant price and marginal cost

The subsection analyses the firm's profit-maximisation problem in a case where products' marginal costs and prices are independent of their attribute levels. That is, the price p_j, the marginal cost c_j, and their differences $(p_j - c_j)$ are fixed for all j. Taking consumer demand $D(x_j;\boldsymbol{x})$ as given, the firm obtains $S(x_j;\boldsymbol{x}|D(x_j;\boldsymbol{x}))$ by maximising profits with respect to q_j^s for each possible $\boldsymbol{x}$: $\max_{q_j^s \in \mathbb{R}_+} \Pi \equiv p \cdot \sum_{j=1}^{n} q_j^D(\boldsymbol{x}) - c \cdot \sum_{j=1}^{n} q_j^s(\boldsymbol{x}) - FC$, subject to $q_j^s(\boldsymbol{x}) \geq q_j^D(\boldsymbol{x})$. The constraint corresponds to the fact that, for any context, quantity supplied of good x_j should be at least as much as its quantity sold – consumers are unable to buy products that have not been produced. With regards to the solution, the first order necessary condition[lx] suggests that the optimal quantity supplied under a context $\boldsymbol{x}$, namely $q_j^{s^*}(\boldsymbol{x})$, equals to quantity sold, $q_j^D(\boldsymbol{x})$, for any option x_j. That is, when market clears, there is no excess supply for any market alternative in the partial equilibrium, i.e., $S(x_j;\boldsymbol{x}|D(x_j;\boldsymbol{x})) = D(x_j;\boldsymbol{x})$.

Considering the learning process, in the short run $q_j^D(\boldsymbol{x})$ may fluctuate substantially since consumers are learning their preferences through experiencing and comparing market alternatives. At this stage, $q_j^{s^*}(\boldsymbol{x})$ and thus $S(x_j;\boldsymbol{x}|D(x_j;\boldsymbol{x}))$[lxi] is not stable. In the longer term, when consumers' preferencesI have been established, $q_j^D(\boldsymbol{x})$ will be fixed at its long-term quantity demanded $q_j^{D^*}(\boldsymbol{x})$ and the firm will produce x_j at $q_j^{s^*}(\boldsymbol{x}) = q_j^{D^*}(\boldsymbol{x})$.

[lx] The Lagrangian of this maximisation problem is $\mathcal{L}(q_j^s, \lambda) = p \cdot \sum_{j=1}^{n} q_j^D(\boldsymbol{x}) - c \cdot \sum_{j=1}^{n} q_j^s(\boldsymbol{x}) - FC + \lambda(q_j^s - q_j^D)$. This gives the Karush-Kuhn-Tucker conditions: (1) $\frac{\partial \mathcal{L}}{\partial q_j^s} = -c + \lambda = 0$, (2) $q_j^s - q_j^D \geq 0$, (3) $\lambda \geq 0$, and (4) $\lambda(q_j^s - q_j^D) = 0$. The solution is $(q_j^{s^*}, \lambda^*) = (q_j^D, c)$.

[lxi] This is because $q_j^{s^*}(\boldsymbol{x})$ is an output of the function $S(x_j;\boldsymbol{x}|D(x_j;\boldsymbol{x}))$ for each $\boldsymbol{x}$.

The firm's second step in solving the profit-maximisation problem is to choose $\boldsymbol{x}^*$ such that $\boldsymbol{x}^* \in \underset{\boldsymbol{x} \subseteq X}{\operatorname{argmax}}\ (p-c) \cdot \sum_{j=1}^{n} D\big(x_j; \boldsymbol{x}\big) - FC$ [lxii]. The first order and second order conditions imply that $(p-c) \cdot \frac{\partial \sum_{j=1}^{n} D(x_j; \boldsymbol{x}^*)}{\partial \boldsymbol{x}^*} = 0$ and $(p-c) \cdot \frac{\partial^2 \sum_{j=1}^{n} D(x_j; \boldsymbol{x}^*)}{\partial (\boldsymbol{x}^*)^2} < 0$, respectively. Specifically, the first order necessary condition suggests that for any $(p-c) \neq 0$, the profit is optimised at a context that either maximises or minimises total demand. The second order condition further indicates that the sign of $(p-c)$ and $\frac{\partial^2 \sum_{j=1}^{n} D(x_j; \boldsymbol{x}^*)}{\partial (\boldsymbol{x}^*)^2}$ is different and neither of them is zero. This means that if $(p-c) > 0$, the optimal context is the one that maximises market demand. Conversely, if $(p-c) < 0$, the optimal context minimises market demand. Accordingly, the possible solution is selecting the context(s) that maximise total market quantity demanded as long as $(p - c) > 0$. Note that this solution implies that $\boldsymbol{x}^*$ is not necessarily chosen to make $\frac{\partial D(x_j; \boldsymbol{x})}{\partial \boldsymbol{x}} = 0$ for all j. That is, the firm only cares whether products are bought, but which one is bought is not of interest. Because of this, the major concern of the firm is whether the context can cause at least one choice pair to pass the threshold criterion during learning [lxiii]. The accuracy of relative rank estimation and the tendency to adopt the rank-order decision rule are far less important in this case. Therefore, the firm has almost no incentive to take advantage of the compromise effect by manipulating menus in a certain direction.

2.1.3.2 Case 2: Attribute-value dependent price and marginal cost

This subsection explores the case where products' marginal costs and prices linearly and positively depend on their attribute values. Consider a case where selling products with higher values of the attribute *m* is more profitable. In this case, the exogenous p_j and

---

[lxii] I note that $D\big(x_j; \boldsymbol{x}\big)$ represents sequences of several discrete data points. To treat it as a differentiable function, it is assumed that $D\big(x_j; \boldsymbol{x}\big)$ had been transformed into a polynomial, continuous function by using Newton's divided differences approach.

[lxiii] The existence of one pair that passes the threshold criterion is sufficient to give rise to a non-flat preferences$_E$ curve.

c_j are no longer constant over all j. Instead, their values depend on the attribute level of x_j on the attribute m. To avoid trade-offs in the price-attribute space[lxiv], assume p and c are a linear function of values of the attribute m such that $\frac{dp(s^m)}{ds^m} > 0$ and $\frac{dc(s^m)}{ds^m} > 0$ for all feasible s^m. For simplicity, let $p_j = \delta_p \cdot s_j^m$ and $c_j = \delta_c \cdot s_j^m$, where $\delta_p, \delta_c \in \mathbb{R}_{++}$.

The first result of the profit-maximising problem in this case is that the firm's optimal quantity supplied under a context $\boldsymbol{x}$ should be equal to the quantity sold for any option x_j, consistent with the previous case. To illustrate, rewrite the total profit earned under a particular context $\boldsymbol{x}$ as $\Pi(\boldsymbol{x}) \equiv \sum_{j=1}^{n} \pi(x_j; \boldsymbol{x}) = \sum_{j=1}^{n} \left[\left(\delta_p \cdot s_j^m\right) \cdot q_j^D(\boldsymbol{x}) - \left(\delta_c \cdot s_j^m\right) \cdot q_j^s(\boldsymbol{x}) \right] - FC \Leftrightarrow \Pi(\boldsymbol{x}) = s_1^m \cdot \left[\delta_p \cdot q_1^D(\boldsymbol{x}) - \delta_c \cdot q_1^s(\boldsymbol{x})\right] + (d + s_1^m) \cdot \left[\delta_p \cdot q_2^D(\boldsymbol{x}) - \delta_c \cdot q_2^s(\boldsymbol{x})\right] + \cdots + [(n-1)d + s_1^m] \cdot [\delta_p \cdot q_n^D(\boldsymbol{x}) - \delta_c \cdot q_n^s(\boldsymbol{x})] - FC$. With knowledge of consumer demand $D(x_j; \boldsymbol{x})$, the firm obtains $S\left(x_j; \boldsymbol{x} \middle| D(x_j; \boldsymbol{x})\right)$ by maximising profits Π with respect to q_j^s for each possible $\boldsymbol{x}$, subject to $q_j^s(\boldsymbol{x}) \geq q_j^D(\boldsymbol{x})$. Not surprisingly, the solution suggests that the profit-maximising output level, $q_j^{s^*}(\boldsymbol{x})$, that constitutes $S\left(x_j; \boldsymbol{x} \middle| D(x_j; \boldsymbol{x})\right)$ is same as previously: $q_j^{s^*}(\boldsymbol{x}) = q_j^D(\boldsymbol{x})$. This implies that no excess supply exists in the market equilibrium.

Likewise, in the short run $q_j^{s^*}(\boldsymbol{x})$ may change considerably with the progress of learning. The supply of products becomes stable in the long run, when consumers have enough experience to choose on a regular basis. As stated before, in the long run, $q_j^{s^*}(\boldsymbol{x})$ equals a constant quantity demanded $q_j^{D^*}(\boldsymbol{x})$.

Secondly, it is found that, as intuition suggests, profit-maximising firms will under plausible assumptions choose the context that maximises the sum of "attribute value" sold. With complete information on $q_j^{s^*}(\boldsymbol{x})$ and $q_j^{D^*}(\boldsymbol{x})$ for all possible $\boldsymbol{x}$, the firm will choose an optimal context $\boldsymbol{x}^*$ so as to maximise its profits. Mathematically, the monopolist will choose $\boldsymbol{x}^* \in \underset{\boldsymbol{x} \subseteq X}{\operatorname{argmax}} \left(\delta_p - \delta_c\right) \cdot \sum_{j=1}^{n} \left[s_j^m(\boldsymbol{x}) \cdot D\left(x_j; \boldsymbol{x}\right)\right] - FC$. By the same token, the

[lxiv] One of the main purposes of this paper is to show that the context effects may be present in the absence of attribute trade-offs.

first and second order conditions imply respectively that $(\delta_p - \delta_c) \cdot \frac{\partial \sum_{j=1}^{n}[s_j^m(\boldsymbol{x}^*) \cdot D(x_j;\boldsymbol{x}^*)]}{\partial \boldsymbol{x}^*} = 0$ and $(\delta_p - \delta_c) \cdot \frac{\partial^2 \sum_{j=1}^{n}[s_j^m(\boldsymbol{x}^*) \cdot D(x_j;\boldsymbol{x}^*)]}{\partial (\boldsymbol{x}^*)^2} < 0$. This reveals that the attribute values of products in this case have a direct effect on the choice of optimal context. The possible solution is to choose the contexts that maximise total attribute values sold as long as $(\delta_p - \delta_c) > 0$.

The solution can be illustrated by a simple example. Rewrite the profit as $\Pi(\boldsymbol{x}) = (\delta_p - \delta_c) \cdot \sum_{j=1}^{n}[s_j^m \cdot q_j^{D^*}(\boldsymbol{x})] - FC$. Suppose that for any two possible contexts $\boldsymbol{x_{old}}$ and $\boldsymbol{x_{new}}$, where $\boldsymbol{s^m_{x_{new}}} = \boldsymbol{s^m_{x_{old}}} + \varepsilon$, $\varepsilon > 0$, the quantities demanded of options with same relative rank under $\boldsymbol{x_{old}}$ and $\boldsymbol{x_{new}}$ are equal (i.e., $q_j^{D^*}(\boldsymbol{x_{old}}) = q_j^{D^*}(\boldsymbol{x_{new}}) \forall j$). Then, the firm's best strategy is to choose the context with higher attribute values, namely $\boldsymbol{x_{new}}$, over $\boldsymbol{x_{old}}$ as long as $\frac{\delta_p}{\delta_c} > 1$ and $\sum_{j=1}^{n} q_j^{D^*} > 0$[lxv]. This example implies that if (by chance) consumer demand for each market alternative x_j is identical across contexts, the firm is better off setting $\boldsymbol{s^m_x}$ as high as possible, including increasing the magnitude of s_1^m or d, as long as the price exceeds marginal costs or at least that one $q_j^{D^*} > 0$[lxvi].

[lxv] Formally, let options in $\boldsymbol{x_{old}}$ and $\boldsymbol{x_{new}}$ be labelled identically as $\{x_1, x_2, x_3, \ldots, x_n\}$, with $\boldsymbol{s^m_{x_{old}}} = \{s_1^m, s_2^m, s_3^m, \ldots, s_n^m\}$ and $\boldsymbol{s^m_{x_{new}}} = \{s_1^m + \varepsilon, s_2^m + \varepsilon, s_3^m + \varepsilon, \ldots, s_n^m + \varepsilon\}$, where $s_{j+1}^m - s_j^m = d$ for all $j = 1, \ldots, n-1$. Express the profits under $\boldsymbol{x_{new}}$ as $\Pi(\boldsymbol{x_{new}}) = (\delta_p - \delta_c) \cdot \sum_{j=1}^{n}[(s_j^m + \varepsilon) \cdot q_j^{D^*}(\boldsymbol{x_{new}})] - FC$. If $q_j^{D^*}(\boldsymbol{x_{old}}) = q_j^{D^*}(\boldsymbol{x_{new}}) = q_j^{D^*}$, the difference between profits will be $\Pi(\boldsymbol{x_{new}}) - \Pi(\boldsymbol{x_{old}}) = (\delta_p - \delta_c) \cdot \sum_{j=1}^{n}[(s_j^m + \varepsilon) \cdot q_j^{D^*} - s_j^m \cdot q_j^{D^*}] = (\delta_p - \delta_c) \cdot \sum_{j=1}^{n}[\varepsilon \cdot q_j^{D^*}] = [(\frac{\delta_p}{\delta_c} - 1) \cdot \varepsilon \cdot \sum_{j=1}^{n} q_j^{D^*}] \cdot \delta_c$. For a monopolist to benefit by choosing $\boldsymbol{x_{new}}$ over $\boldsymbol{x_{old}}$, the condition that $\Pi(\boldsymbol{x_{new}}) - \Pi(\boldsymbol{x_{old}}) > 0$ is required. As, by definition, $\delta_c > 0$, $\Pi(\boldsymbol{x_{new}}) - \Pi(\boldsymbol{x_{old}}) > 0 \Leftrightarrow (\frac{\delta_p}{\delta_c} - 1) \cdot \varepsilon \cdot \sum_{j=1}^{n} q_j^{D^*} > 0$. This implies that $\frac{\delta_p}{\delta_c} > 1$ and $\sum_{j=1}^{n} q_j^{D^*} > 0$ as ε is strictly positive by definition.

[lxvi] Importantly, even though these two conditions are satisfied, the firm may choose to producing nothing if FC is too high such that $0 > \Pi(\boldsymbol{x_{new}}) > \Pi(\boldsymbol{x_{old}})$.

Consider another special situation where consumers are highly "extremeness averse", so that they always purchase the middle option, say x_2, in a three-option set regardless of its actual attribute level. The firm's optimal context choice here is to produce x_2 only, and to set s_2^m as large as possible as long as $q_2^{D^*}(\boldsymbol{x}) > 0$[lxvii]. In other words, if the quantity demanded of the middle option is positive, the firm can enjoy an increase in profits by shifting $\boldsymbol{s_x^m}$ towards its upper limit. Together, consistent with the solution, these two examples illustrate the fact that when levels of products' common attributes are directly associated with profits, the firm has a monetary incentive to manipulate the product line to increase the quantity demanded for more profitable products.

However, consumer behaviour in the proposed model is more complex than in these two special examples. In effect, once s_j^m exceeds the value that produces the maximum preference$_U$, any increment in s_j^m will incur a loss in q_j^D, which reflects a trade-off in determining the menu. Yet, $\boldsymbol{x}^*$ here may be closer to the upper limit of the attribute space than in the first class of the profit-maximisation problem, since the presence of s_j^m in the first and second order condition allows profit-maximising context to (slightly) move away from the preference$_U$-maximising point. Moreover, in this case, consumers who suffer a stronger degree of social influence may be exploited more than those who are relatively less biased. This is because the false consensus effect may mislead biased consumers to choose the middle option when others are actually preferred. As a result, it is predicted that the firm will be incentivised to distort context information to take advantage of consumer bias in this case.

2.1.3.3 Summary

This formulation of the model suggests that when marginal costs and prices are constant, the firm is best off producing a context that maximises total market quantity demanded, as long as prices exceed marginal costs. That is, the distribution of market demand over market alternatives does not affect profit. Therefore, there is no incentive

---

[lxvii] This result can be obtained by plugging in $q_1^{D^*} = q_3^{D^*} = 0$ into the previously-derived condition $\sum_{j=1}^{n} q_j^{D^*} > 0$.

for the monopolist to exploit biased consumers by distorting context. In contrast, results from the second case, where products' marginal costs and prices are linearly dependent on attribute values, suggest that the attribute values of products have a direct effect on the choice of context. The firm here is best off choosing contexts that maximise total attribute values sold. The firm hence has a monetary incentive to manipulate the product line to increase the quantity demanded for more profitable products.

2.3 Computer simulation of the compromise effect

2.3.1 Preliminaries

This section illustrates the model's predictions about how market context and preferences$_U$ affect consumer decisions made from experience. Focussing on the compromise effect, the model argues that context-sensitive behaviour does not necessarily arise from violations of rationality assumption[lxviii], lack of market experience, or difficulty in trade-offs among common attributes. In fact, biased market inference *per se* can be a contributor to the existence of the compromise effect in long-term equilibrium. To demonstrate this, the following sections will use several simulations to illustrate consumer decision-making and derive the boundary conditions of the compromise effect, the firm's optimal responses, and consumer welfare loss in the partial equilibrium. The simulation will first use a single agent to explain the model's mechanism in more details and then use multi-agents to simulate real market scenario. The proposed model is programmed with Matlab 2020b.

2.3.1.1 A rational choice benchmark model

To investigate how consumer choices predicted by the proposed model deviate from theoretically optimal choices, a benchmark model that operates within the general framework of rational choice theories will be presented in the simulation. In line with the

[lxviii] This includes a violation of any standard axiom of preferences and a violation of value/ utility maximisation.

homo economicus assumptions of rationality, the benchmark consumer *i* is assumed to physically evaluate all market offerings to acquire a subjective valuation of each of them, i.e., to know the absolute magnitude of $PU_i^m(x_j)$, from sensory experiences[lxix], and then rationally select the one with highest PU_i^m. To reduce the complexity of comparisons of predictions, the threshold criterion used in the benchmark model implies that the consumer *i* will choose nothing or an outside option, x_0, if the difference between the preference$_U$ of the most preferred available option and the least preferred available option is smaller than the threshold t[lxx]. This is because as long as the PU_i^m of the most preferred option differs from the least preferred one by an amount greater than *t*, a pairwise choice will be made. Therefore, even if it is the only pair that passes the threshold criterion, eventually the consumer *i* will still buy a product in the market. Moreover, the value of *t* is the same for a given consumer no matter which model is used since it is assumed to be a fixed dispositional trait. Given a non-binding budget constraint, once the threshold criterion is passed, the unit-demand consumer *i* will choose the option that maximises PU_i^m under the current context. If there is more than one option that maximises it, one of

---

[lxix] Since the exact attribute levels of products are here assumed not to be shown (or are shown but not attended to) on the products' labels, consumers have to experience the products to know their subjective values.

[lxx] One can justify this assumption by viewing the threshold, *t*, as an opportunity cost of buying in the market. The rational agent buys the product only if its subjective value exceeds *t* (i.e., the threshold criterion is a necessary condition for a product being brought). The use of relative comparison ($PU_i^m\left(x_j\right) - PU_i^m(x_k)$ for $x_j \neq x_k \wedge \ x_j, x_k \in \boldsymbol{x}$) is to make this condition consistent with the one in the proposed model, so that the effects of preference uncertainty and biased learning can be detected more easily by comparisons. This can be justified by the following intuition. The consumer may feel it is painful to forgo any option with high PU_i^m. If the valuations of all options are high and arbitrarily close to each other, consumers may choose an outside option to avoid making decisions among them, even though they all are good in an absolute sense. Consumers may be more willing to choose among available products if there exists at least one option associated with a relatively low PU_i^m.

these options will be chosen randomly. In addition, past contexts will not affect consumer decision making in the benchmark model.

2.3.1.2 A measure of the compromise effect

For all simulations conducted in this section, the product line is assumed to comprise five market offerings ordered along one dimension. Therefore, the analysis of contextual influences will be based on quinary-quinary set comparisons. As suggested by Simonson and Nowlis (2000), one way to measure the compromise effect is called the "middle proportions" approach[lxxi], which treats the probability of choosing the middle option as an indicator of the propensity to compromise. However, a relatively high choice probability of the middle option may result from its high associated preferences$_U$. That is, the middle option itself may simply be more intrinsically desirable than any other choice alternatives. A high tendency to choose the middle option in this case is irrelevant to the chosen option's rank in the choice set, making it inconsistent with the definition of the

---

[lxxi] Another two frequently employed measures are changes in absolute share (Simonson, 1989) and changes in relative share (Simonson & Tversky, 1992). The former compares the market share of an option before and after it becomes the middle one, and the latter assesses changes in relative share of the same choice pair, of which one option should be at the middle of the choice set, before and after context shifts. The reason why these two measurement paradigms are not used in this section is mainly because they are not suitable for analysing the case where choice probability is either 0 or 1. For example, suppose the predicted quantity demanded for products x_1, x_2, x_3, x_4 is $(q_1^{D^*}, q_2^{D^*}, q_3^{D^*})$ = (0, 0, 1) in context 1 and $(q_2^{D^*}, q_3^{D^*}, q_4^{D^*})$ = (0, 1, 0) in context 2. The changes in absolute/ relative share approach may conclude there is no such compromise effect since $q_3^{D^*}$ is the same across contexts and the relative choice probability of x_2 and x_3 is the same. Suppose the benchmark quantity demanded is $(q_1^{D^*}, q_2^{D^*}, q_3^{D^*})$ = (0, 0, 1) in context 1 and $(q_2^{D^*}, q_3^{D^*}, q_4^{D^*})$ = (0, 0, 1) in context 2. The changes in absolute and relative share of x_3 should be bounded between 1 and 0. If the model predictions are compared with benchmark results, the deviation may be detected. However, it is difficult and complex to implement this kind of comparison quantitatively. Therefore, I decided to drop these two paradigms.

compromise effect. To address this issue, in this all following simulations, the compromise effect will be assessed by comparing the model's estimated probability of choosing middle options with the benchmark model's probability under each choice set. A comparison with benchmark results allows the estimated strength of the compromise effect to rule out the possibility that an observed concentration of choice on the middle option is purely thanks to the middle option's high associated preference$_U$, relative to other choice alternatives'. Moreover, the degree of the compromise effect may vary with context. By comparing against the benchmark, contextual influences on consumer choices can be understood in relation to other choice drivers such as actual attribute values of products, because the benchmark does not capture the impact of the relative rank position within a choice set.

2.3.1.3 Consumer welfare

Consumption choices will be appraised from a hedonic perspective in order to estimate consumer detriment. Suppose that the output of the preference$_U$'s representative function is quantified in a cardinal term. The consumer i's psychological welfare loss[lxxii] under each market context will be computed by subtracting the (averaged) PU_i^m of the predicted choice(s) from the benchmark. In addition, since the firm may take advantage of consumer bias[lxxiii], welfare loss resulting from the firm's manipulation of context will be examined by comparing the consumer welfare in the present model's predicted equilibrium with the welfare level of the choice predicted from the proposed model under the equilibrium context in the benchmark model. The welfare losses in these two equilibria will also be compared.

[lxxii] Despite consumers being potentially unaware of their loss of happiness, it is important to investigate potential negative effects of contextual influences on consumer welfare.

[lxxiii] Recall that for some product categories, it is more profitable to sell products with higher attribute values. If consumers show a bias towards the middle option, the monopolist can be benefited from distorting the product line to increase the demand for its more profitable products.

2.3.2 Simulation 1: The compromise effect

To illustrate the compromise effect and its potential influences on partial market equilibrium, consider a situation with a single agent (or multiple identical agents) and two menus $\boldsymbol{x_1}$ = {x_1, x_2, x_3, x_4, x_5} and $\boldsymbol{x_2}$ = {x_2, x_3, x_4, x_5, x_6}[lxxiv], where all alternatives are arranged in order according to their values of the attribute *m*. In this simulation, let $\boldsymbol{s^m_{x_1}}$ = {0.25, 0.35, 0.45, 0.55, 0.65} and $\boldsymbol{s^m_{x_2}}$ = {0.35, 0.45, 0.55, 0.65, 0.75}.

<u>2.3.2.1 The benchmark model</u>

Suppose that the representative consumer *i* is endowed with a context-independent preference$_U$, PU_i^m = Beta (5, 15) and a threshold parameter *t* = 0.5. As shown in Figure 2 (a), the benchmark choices under $\boldsymbol{x_1}$ and $\boldsymbol{x_2}$ are x_1 and x_2, respectively, since $PU_i^m(s_1^m) - PU_i^m(s_j^m) > 0.5\ \forall\ j = 2, 3, 4, 5$ and $PU_i^m(s_2^m) - PU_i^m(s_k^m) > 0.5\ \forall\ k = 3, 4, 5, 6$. Thus, the revealed preference orderings here do not change with context.

<u>2.3.2.2 Choice estimated by the proposed model and the compromise effect</u>

If the consumer *i* cannot consciously access the absolute magnitude of preferences$_U$, the choices may be, to some degree, governed by the self-perceived relative position in the taste distribution and the tendency to leverage context-generated information. This subsection assesses choice behaviour of three types of consumer *i* under each context using the proposed model and then compare their tendency to choose the middle option from a menu with the benchmark results. The types of consumers investigated are as follows: (1) a type that does not trust any preference-related information conveyed in the menu at all, i.e., Φ = 0; (2) a type that assigns equal attention to internal and external cues during learning, i.e., Φ = 0.5, and suffers some degree of the false consensus effect, and (3) a type that pays attention only to external cues, i.e., Φ = 1, and is susceptible to the false consensus effect like the second type but to a greater extent.

[lxxiv] In terms of relative ranks, options would be labelled as x_1, x_2, x_3, x_4, x_5 under both contexts. The use of $\boldsymbol{x_2}$ = {x_2, x_3, x_4, x_5, x_6} is to emphasise that the *j*th option in $\boldsymbol{x_1}$ is the (*j*-1)th option in $\boldsymbol{x_2}$, with respect to their attribute values. This is just to make the explanation of the compromise effect more clear.

Figure 2 (b) shows the preference$_U$, the preference$_N$, and the preference$_C$ for the Φ = 0.5 type, under $\boldsymbol{x_1}$ and $\boldsymbol{x_2}$. Notice that during every pairwise comparison, if a pair passes the threshold criterion, the preference that guides the pairwise choice will be the preference$_U$ when $\Phi = 0$, the preference$_C$ when $\Phi = 0.5$, and the preference$_N$ when $\Phi = 1$. The record of choice probability of each option will then be translated into the preference$_I$, the preference that informs final choice. The preferences$_I$ for $\Phi = \{0, 0.5, 1\}$ are shown in Figure 2 (c) and (d).

As shown in these figures, the choice made by the consumer with $\Phi = 0$ is identical to the benchmark one as they have the same threshold criterion and strictly adhere to their preference$_U$. Interestingly, the fact that the preference$_I$ curves of the $\Phi = 0$ and 0.5 types are the same suggests that, under the current parameter values, the bias in the social norm perception is not strong enough to effectively alter the final choice when the weight assigned to the external cues is moderate. This is because in $\boldsymbol{x_1}$, the dominant option, x_1, substantially outstrips the less extreme options like x_2 and x_3. As a result, the bonus gained from the preference$_N$ is not sufficient to offset the inferiority in the preference$_U$ (see Figure 2 (b)). Yet, if the consumer *i* can completely ignore the internal cues after deciding to choose from the pair, i.e., if $\Phi = 1$, the preference$_N$ will take full control in the second stage of the pairwise comparison. This significantly increases the probability of choosing x_2 and x_3 from pairs and leads the consumer *i* with $\Phi = 1$ to eventually act as if x_2 is mostly preferred in $\boldsymbol{x_1}$, as shown in the yellow preference$_I$ line in Figure 2 (c).

Note that in the case of $\Phi > 0$, the preferences$_N$ come to affect pairwise choice only at the second stage of comparison, which occurs only when a difference in the preferences$_U$ between pair alternatives exceeds a threshold value. This means that whether the preferences$_N$ (or Φ) have an opportunity to influence choice and the preferences$_I$ is solely determined by preferences$_U$. The indispensability of the preferences$_U$ in decision making explains why the middle-ranking option is not chosen most often in $\boldsymbol{x_1}$ even when Φ =1.

To explore this explanation in more detail, consider the following example. Suppose there is a middle option, say x_{Mid}, whose associated preference$_U$ value is fairly equal to the threshold value. There are three possible situations regarding pairwise comparisons that x_{Mid} is involved in. First, if x_{Mid} is paired with an option that is associated with an even lower preference$_U$ value, it is highly possible that the pair cannot pass the threshold criterion due to a low preference$_U$. In this situation, the preferences$_N$ play no role in pairwise choice and hence in the formation of preferences$_I$. Second, if x_{Mid} is paired with an option whose associated preference$_U$ value is slightly higher than x_{Mid}'s, the threshold criterion still cannot be passed. The result is same as in the previous case. Third, if x_{Mid}'s paired option is associated with a much higher preference$_U$ value, the threshold criterion is highly likely to be passed. The preferences$_N$, finally, can have an effect on pairwise choice and the preferences$_I$, with the degree of the effect depending on preferences$_U$, preferences$_N$, and Φ. When Φ is large, e.g., 1, x_{Mid} can of course receive more choice share than its paired option in this case. However, this pair will not be the only pair that pass the threshold criterion. The option that has a high preference$_U$ value can still generate choice share from other pairwise comparisons. In the end, the advantage that Φ and preferences$_N$ give to x_{Mid} may not be enough for it to have the highest accumulated choice share gained from pairwise comparisons among all choice alternatives. This is exactly what happens to x_3 in $\boldsymbol{x_1}$ (and even $\boldsymbol{x_2}$) when Φ = 0.5 and 1.

I return now to the discussion of consumer choice estimated by the proposed model. In $\boldsymbol{x_2}$, the far higher preference$_U$ value still gives the first option, x_2, an unsurpassable advantage in pairwise comparison when the interference of external forces is moderate (Φ = 0.5). Therefore, the preference$_C$ in $\boldsymbol{x_2}$ shown in Figure 2 (b) is somewhat downward-sloping, similar to the pattern of the preference$_U$ curve. This then causes consumers with Φ = 0 and Φ = 0.5 to exhibit identical preference$_I$ curves. Conversely, if Φ is large enough to significantly amplify the relative influence of the preference$_N$ on the pairwise choice, then the advantage of associating with the highest preference$_U$ value may be overcome. Figure 2 (d) shows that when the second stage of the pairwise comparison is totally controlled by the preference$_N$, i.e., Φ = 1, the second option, x_3, will instead be chosen more frequently over other options in $\boldsymbol{x_2}$. Apparently, the prediction that the consumer with Φ = 1 chooses the second option under both $\boldsymbol{x_1}$ and $\boldsymbol{x_2}$ implies that x_2 is preferred to x_3 under $\boldsymbol{x_1}$, whereas the option x_3 is preferred to x_2 under $\boldsymbol{x_2}$, indicating a preference reversal. Yet, in fact, the consumer i does not make an optimal decision under each of contexts. Compared with the benchmark consumer, the consumer with Φ = 1 exhibits a strong propensity to choose the second option under both $\boldsymbol{x_1}$ and $\boldsymbol{x_2}$. This shows that even given a stable preference$_U$ and market experience, uncertainty over the preference$_U$ and the false consensus effect may bias inferences about the preference$_U$. Consequently, the consumer acts as if the relative ordering of x_2 and x_3 is not consistent.

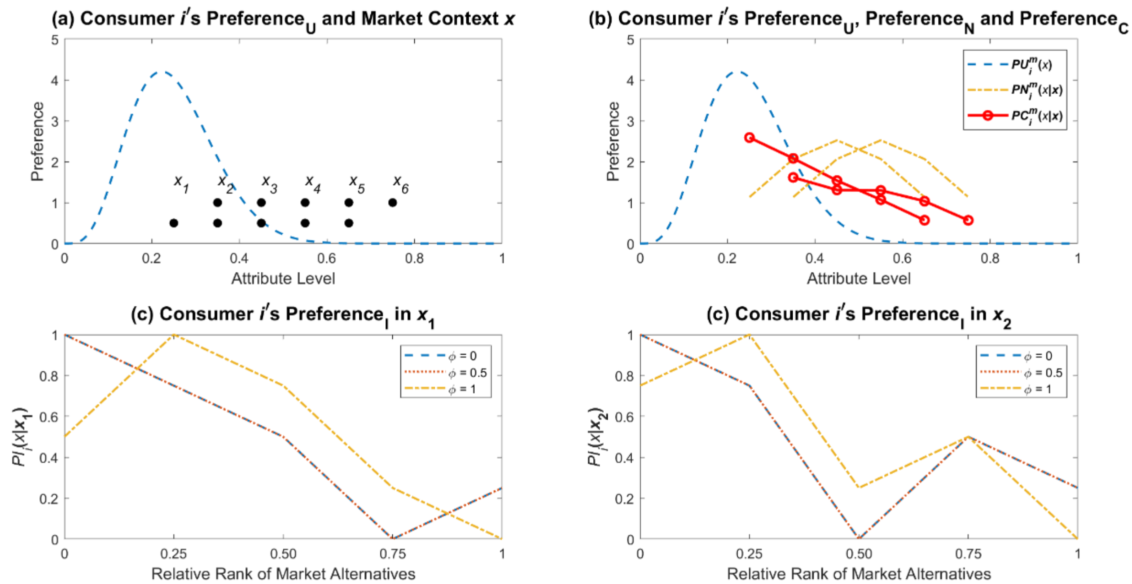

Figure 2. Simulation results for the demand side. (a) The consumer i's preference$_U$ and market offerings; (b) Consumer preferences under each context on one scale. The bell-shaped preference$_N$ reflects the existence of the false consensus effect; (c) The rank-based preference$_I$ in the first context. The zero preference$_I$ value of the fourth option possibly results from the fact that $PU_i^m(x_4)$ is too high to make pairs (x_3, x_4) and (x_4, x_3) pass the threshold criterion, whereas x_5 can obtain some pairwise choice probability from pairs (x_3, x_5) and (x_5, x_3). Notice that this does not happen in the case where Φ = 1 since the strong tendency to use biased market inference effectively raises the choice probability of x_4 in other pairs; (d) The preference$_I$ in the second context. The reason why the preference$_I$ reaches zero or substantially slumps at the middle is similar to the one stated in (c).

2.3.2.3 The monopolist's optimal strategy, equilibrium contexts, and consumer welfare: Case 1

This subsection explores the firm's profit-maximising strategy, given market demand and equilibrium context in the case of a fixed price and marginal cost. Consumer welfare under market equilibrium context is also examined. Consider the case where the price p and the marginal cost c are constant for all market alternatives. Given what the unit-demand consumer i will choose in the market under both contexts, the firm is indifferent between providing $\boldsymbol{x_1}$ and $\boldsymbol{x_2}$. If $(p - c) \cdot q^{D^*} \geq FC$, the partial equilibrium context will be $\boldsymbol{x}^* = \{\boldsymbol{x_1}, \boldsymbol{x_2}\}$, with the consumer in the proposed model choosing x_2 in $\boldsymbol{x_1}$ and x_3 in $\boldsymbol{x_2}$ when $\Phi = 1$ and choosing x_1 in $\boldsymbol{x_1}$ and x_2 in $\boldsymbol{x_2}$ when $\Phi = 0$ and 0.5. Otherwise, the firm may be better off producing nothing. Notice that for all three types, the consumer i is better off if the firm produces $\boldsymbol{x_1}$ rather than $\boldsymbol{x_2}$, since consumer welfare obtaining from consumption is higher under $\boldsymbol{x_1}$ than under $\boldsymbol{x_2}$[lxxv]. Under $\boldsymbol{x_1}$, welfare loss resulting from biased learning is larger than the welfare loss under $\boldsymbol{x_2}$[lxxvi].

2.3.2.4 The monopolist's optimal strategy, equilibrium contexts, and consumer welfare: Case 2

The second case examines the scenario where the prices and marginal costs of products are linearly associated with their attribute levels. The firm's optimal responses to market demand, equilibrium contexts, and consumer welfare under each context are investigated in this subsection.

[lxxv] In the benchmark model and the $\Phi = \{0, 0.5\}$ cases, the consumer i gains $PU_i^m(x_1) =$ 4.0466 by choosing x_1 under $\boldsymbol{x_1}$ and gains $PU_i^m(x_2) = 2.0967$ by choosing x_2 under $\boldsymbol{x_2}$. The fact that $PU_i^m(x_1) > PU_i^m(x_2)$ suggests that the consumer i is better off under $\boldsymbol{x_1}$. Likewise, $PU_i^m(x_2) > PU_i^m(x_3) = 0.5526$ suggests that in the $\Phi = 1$ case, the consumer i will be better off under $\boldsymbol{x_1}$.

[lxxvi] Under $\boldsymbol{x_1}$, the welfare loss due to compromise behaviour is computed as $PU_i^m(x_1) - PU_i^m(x_2) = 1.9499$, and the welfare loss under $\boldsymbol{x_2}$ is $PU_i^m(x_2) - PU_i^m(x_3) = 1.5441$, which is smaller than the former.

Assume that $c_j = \delta_c \cdot s_j^m$ and $p_j = \delta_p \cdot s_j^m$, and the profit function is $\Pi(\boldsymbol{x}) = (\delta_p - \delta_c) \cdot \sum_j [s_j^m \cdot q_j^{D^*}(\boldsymbol{x})] - FC$. Given consumer choice predicted above, for all cases of Φ and the benchmark, the monopolist has an incentive to set the menu to be $\boldsymbol{x_2}$ rather than $\boldsymbol{x_1}$ as long as $\frac{\delta_p}{\delta_c} > 1$[lxxvii]. This implies that when q^{D^*} is high, enlarging the gap between δ_p and δ_c generates more monetary benefits of changing context from $\boldsymbol{x_1}$ to $\boldsymbol{x_2}$ than when q^{D^*} is low, provided that $\frac{\delta_p}{\delta_c} > 1$.

In a nutshell, if $\frac{\delta_p}{\delta_c} > 1$ holds and $\Pi(\boldsymbol{x_2}) \geq 0$, the partial equilibrium context will be $\boldsymbol{x}^* = \boldsymbol{x_2}$ and the consumer with $\Phi = 1$ chooses x_3 and those with $\Phi = 0$ and 0.5 chooses x_2. If any condition does not hold, the firm will be better off by producing $\boldsymbol{x_1}$, or choosing to shut down, depending on whether $\Pi(\boldsymbol{x_1}) \geq 0$. As mentioned, consumer welfare is lower at $\boldsymbol{x_2}$ than $\boldsymbol{x_1}$. However, this does not mean that the firm distorts the informational content of its product line with the intention of exploiting consumers, since the firm's optimal strategy is same as the one in the benchmark model. In other words, consumers' psychological and monetary loss (paying p_3 instead of p_2) in this case stems from their own biases, not from the firm's manipulation.

2.3.3 Simulation 2: Boundary conditions of the compromise effect

To derive the boundary conditions of the compromise effect, Simulation 1 was extended to the world where the monopolist changes the context each time by raising attribute levels of all current options by 0.01 units till the last option's attribute value equals to 1. The simulation started at $\boldsymbol{s_x^m}$ = {0.01, 0.11, 0.21, 0.31, 0.41} and ended when

---

[lxxvii] It is a special case when the equation $\Pi(\boldsymbol{x_2}) - \Pi(\boldsymbol{x_1}) = [(\frac{\delta_p}{\delta_c} - 1) \cdot 0.1 \cdot q^{D^*}] \cdot \delta_c$ holds for all cases, i.e., Φ = {0, 0.5, 1} and the benchmark. This is because the chosen option in $\boldsymbol{x_1}$ and $\boldsymbol{x_2}$ has the same relative rank and $\boldsymbol{s_{x_2}^m} = \boldsymbol{s_{x_1}^m} + d$. As discussed in the previous section, in this case, the differences between profits do not depend on which option was chosen, but on the total quantity demanded of the chosen option. Importantly, this coincidence is unlikely to occur when there exist many heterogenous consumers since it requires $q_j^{D^*}(\boldsymbol{x_1}) = q_j^{D^*}(\boldsymbol{x_2})\ \forall\ j$.

it reached $\boldsymbol{s_x^m}$ = {0.60, 0.70, 0.80, 0.90, 1.00}. In total there are $\frac{0.6-0.01}{0.01} + 1 = 60$ contexts. Moreover, retaining the same settings as in Simulation 1, consumer i is endowed with PU_i^m = Beta (5, 15), a decision-making threshold of 0.5, and three possible values for weighting parameter: Φ = {0. 0.5, 1}. The preference$_U$ curve and market offerings of contexts $\boldsymbol{x_1}$ and $\boldsymbol{x_{60}}$ are presented in Figure 3 (a).

2.3.3.1 The benchmark model

The simulation first explores choice made by the benchmark consumers under each context. Given that the PU_i^m peaks at $s^m \cong$ 0.2222, the fourth and the fifth options will never be the most preferred ones as they are always dominated by the third option in the set. Moreover, the middle option beats the other four options when context lies between $\boldsymbol{x_1}$ and $\boldsymbol{x_7}$ ($\boldsymbol{s_7^m}$ = {0.07, 0.17, 0.27, 0.37, 0.47}). The benchmark agent, therefore, chooses x_3 from the menu $\boldsymbol{x}$ = {$\boldsymbol{x_1}$, $\boldsymbol{x_2}$, …, $\boldsymbol{x_7}$}, as shown in Figure 3 (b). However, when the context moves to $\boldsymbol{x_8}$, the second option becomes the optimal one since $s_{x_2|\boldsymbol{x_8}}^m$ = 0.18 is closer to 0.2222 than $s_{x_3|\boldsymbol{x_8}}^m$ = 0.28 is. The benchmark agent will then continue choosing x_2 until the context reaches $\boldsymbol{x_{17}}$, where $s_{x_2|\boldsymbol{x_{17}}}^m$ = 0.27 = $s_{x_3|\boldsymbol{x_7}}^m$. Repeating what had happened previously, at $\boldsymbol{x_{18}}$, the dominant position of x_2 is replaced by x_1 as $s_{x_2|\boldsymbol{x_{18}}}^m$ is now 0.28 and $s_{x_1|\boldsymbol{x_{18}}}^m$ becomes 0.18, which is closer to 0.2222. Afterwards, when context goes beyond $\boldsymbol{x_{45}}$, the benchmark consumer chooses nothing as none of the pair differences can pass the threshold criterion. Together, owing to the awareness of the exact value of $PU_i^m(x|\boldsymbol{x})$, choices made by the benchmark consumer perfectly reflect the preference orderings of the market alternatives. Consumer choices are consistent and stable across contexts.

2.3.3.2 Market share estimated using different values of $\boldsymbol{\Phi}$ and the compromise effect

The choices of consumers who are not directly influenced by the magnitudes of preferences$_U$ may be influenced by available options' relative ranks in addition to their absolute attribute levels. This subsection assesses choices made by cognitively-constrained consumers who cannot consciously access the absolute value of their preferences$_U$ using the proposed model. As in Simulation 1, three types of consumers (Φ = {0. 0.5, 1}) are investigated. By comparing their choice behaviour to each other and to the benchmark results, the subsection shows why deliberately using market information and information about social norms to assist decision making may bias decisions.

Firstly, the behaviour of consumers with $\Phi = 0$, mirrors the benchmark result, as indicated in Figure 3 (c). This suggests that unbiased preference learning may bring about optimal outcomes, despite uncertainty about the magnitudes of preferences$_U$.

Conversely, Figure 3 (d) and (e) illustrate that consumers with $\Phi = 0.5$ and 1 are more likely to buy the options close to the middle, compared to the benchmark. More specifically, the choices of consumers who suffer moderate degree of social influences may remain at x_3 or x_2 for a few more contexts when they are no longer the best option in the set. Yet, choice will soon be switched to the optimal one when these previously-dominating options become far worse, e.g., at $\boldsymbol{x_9}$ and $\boldsymbol{x_{21}}$. The compromise effect as measured by the middle proportion approach occurs at context $\boldsymbol{x} = \{\boldsymbol{x_8}, \boldsymbol{x_{18}}, \boldsymbol{x_{19}}, \boldsymbol{x_{20}}\}$. Apart from this, choice reversal is observed when the context is $\{\boldsymbol{x_9}, \boldsymbol{x_{19}}\}$ and $\{\boldsymbol{x_{10}}, \boldsymbol{x_{20}}\}$. For both pairs of contexts, choices in the latter context, i.e., $\boldsymbol{x_{19}}$ and $\boldsymbol{x_{20}}$, imply that x_2 is preferred to x_1, which, however, is not consistent with the preference relation implied in the former context since x_2 (x_1 in the latter context) is chosen over x_3 (x_2 in the latter context) in $\boldsymbol{x_9}$ and $\boldsymbol{x_{10}}$.

The most socially influenced consumers, i.e., where ($\Phi = 1$), choose the middle option until $\boldsymbol{x_{24}}$, with a temporary shift to x_4 and x_2 at $\{\boldsymbol{x_7}, \boldsymbol{x_8}\}$[lxxviii] and $\{\boldsymbol{x_{13}}, \boldsymbol{x_{14}}\}$, respectively. According to Figure 3(e), consumers with $\Phi = 1$ turn back to the optimal option after context moves to $\boldsymbol{x_{36}}$. In this case, the compromise effect defined by the middle proportion measure is seen when the context is context $\boldsymbol{x} = \{\boldsymbol{x_9}, \boldsymbol{x_{10}}, \boldsymbol{x_{11}}, \boldsymbol{x_{12}}\} \cup \{\boldsymbol{x_{15}}, \boldsymbol{x_{16}}, \ldots, \boldsymbol{x_{24}}\}$. Compared with the result obtained in the case of $\Phi = 0.5$, it is suggested that

---

[lxxviii] This is because at these two contexts, the difference between the preference$_U$ of the optimal option and the second-best option is too small to pass the threshold criterion, rendering the pairs involving these two options less likely to be chosen during the pairwise comparison. This largely benefits the third best option, x_4. This does not happen when $\Phi = 0$ and 0.5 possibly because the value of the fourth option's preference$_U$ is too low to be offset by this benefit. The same logic can be applied to the observation at $\{\boldsymbol{x_{13}}, \boldsymbol{x_{14}}\}$, where PU_i^m values of the first and third option are too close.

consumers with Φ = 1 are more likely to exhibit a biased tendency towards the middle option because they trust their biased estimates of their relative positions in population more. Moreover, as in the previous case, choice reversals appear in many pairs of contexts: $\{\boldsymbol{x}_1, \boldsymbol{x}_{11}\}$, $\{\boldsymbol{x}_2, \boldsymbol{x}_{12}\}$, $\{\boldsymbol{x}_5, \boldsymbol{x}_{15}\}$, $\{\boldsymbol{x}_6, \boldsymbol{x}_{16}\}$, $\{\boldsymbol{x}_9, \boldsymbol{x}_{19}\}$, $\{\boldsymbol{x}_{10}, \boldsymbol{x}_{20}\}$, $\{\boldsymbol{x}_{11}, \boldsymbol{x}_{21}\}$, $\{\boldsymbol{x}_{12}, \boldsymbol{x}_{22}\}$, and $\{\boldsymbol{x}_{25}, \boldsymbol{x}_{35}\}$. Apparently, if preferences are legitimately inferred from observations of final choices, one easily concludes that the changes of market context result in preference reversals. However, the relative ranks of available products affect their relative attractiveness without changing preferences$_U$.

In summary, the presence of the compromise effect implies that when the PU_i^m of the middle option is high enough, the false consensus effect will play a prominent role in decision making by biasing individual's preferences$_C$, which inclines consumers towards choosing the middle option. Yet, this also means that consumer choice is not solely based on products' relative rank. If the middle option is not attractive enough in terms of its absolute attribute value, the existence of the context-independent preferences$_U$ and opportunities to experience prevent consumers from exhibiting the compromise behaviour. In other words, the limit of contextual influences is also determined by preferences$_U$ and the market context.

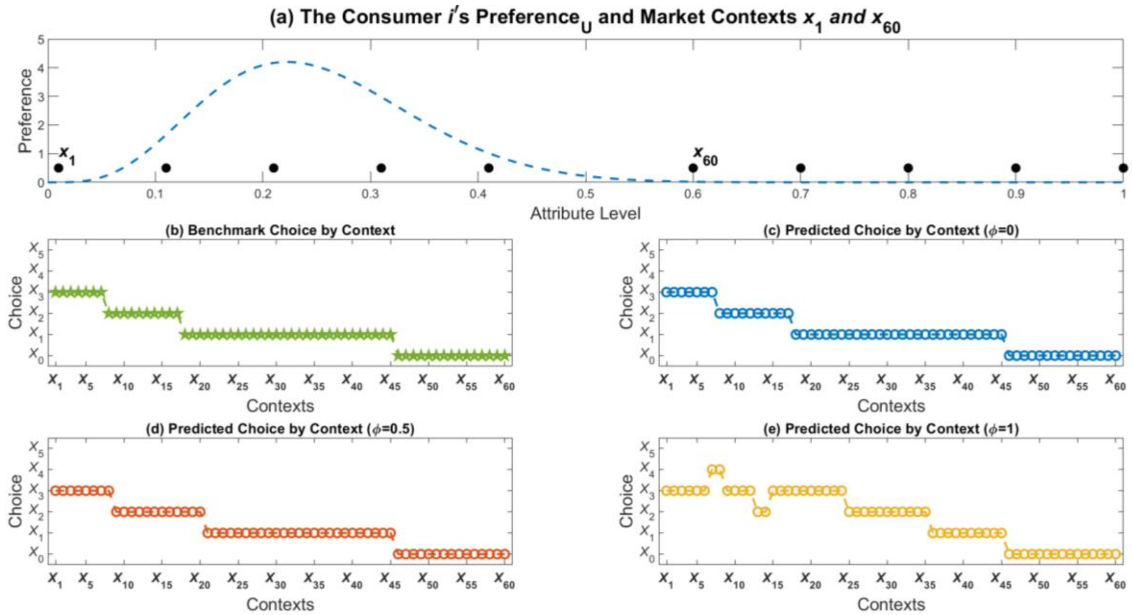

Figure 3. Consumer choice under each feasible context in the respective model. Note that consumers are assumed to buy an outside option, x_0, if they do not make a choice from the market offerings. Therefore, the choice of x_0 shown in figures (b), (c), (d), and (e) indicates a situation where consumers choose to opt out.

2.3.3.3 The monopolist's optimal strategy, equilibrium contexts, and consumer welfare: Case 1

This subsection investigates the monopolist's profit-maximising strategy and equilibrium contexts in the situation where all products have the same price and marginal cost, which is a constant. Consumer welfare for different types of consumers under their respective market equilibrium is also explored.

Firstly, all types of the consumer i consistently choose only one option from the menu from $\boldsymbol{x_1}$ to $\boldsymbol{x_{45}}$, and so the profits gained from setting $\boldsymbol{x}$ within this context range is the same when prices and marginal costs are constant across all feasible products. Mathematically, the profit will be $\Pi = (p - c) \cdot q^{D^*} - FC$. In line with the conclusions of the theoretical analysis and Simulation 1, profits in this case do not depend on which option is chosen, but on the relative magnitudes of exogenously determined prices and marginal costs as well as the total quantity demanded of the chosen option, q^{D^*}. Therefore, the monopolist is indifferent between these contexts. The partial market equilibria are $\boldsymbol{x}^*$ = {$\boldsymbol{x}_1$, $\boldsymbol{x}_5$, …, $\boldsymbol{x}_{45}$} if profits earned under these contexts are positive. Otherwise, the firm will choose to produce nothing.

However, consumer welfare under these equilibria differs substantially. As shown in Figure 4, except for those with Φ = 1, the consumer i enjoys the highest happiness level when $\boldsymbol{x}$ = {$\boldsymbol{x}_2$, $\boldsymbol{x}_{12}$, $\boldsymbol{x}_{22}$} as the attribute value of the chosen option is closest to 0.2222, while gaining the least pleasure from consumption when $\boldsymbol{x}$ = $\boldsymbol{x}_{45}$. Additionally, under contexts where the compromise effect exists, consumer experience welfare loss to varying degrees because of cognitive constraints and biased social sampling. Obviously, consumers with Φ = 1 suffer substantially more than other two types of consumers do, as shown in Figure 4. If the society consists of equal amounts of different types of consumers, the socially optimal equilibrium context in this case will be $\boldsymbol{x}_2$ where both biased and unbiased consumers make an optimal choice.

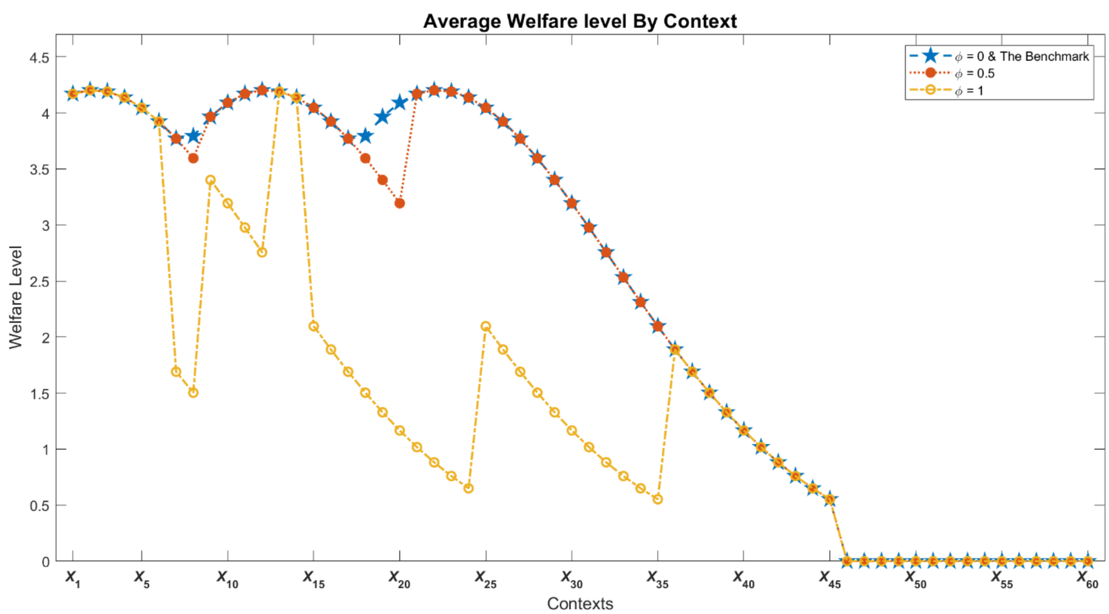

Figure 4. The average welfare level resulting from buying in the market. The benchmark welfare level is identical to the one that the (Φ = 0) type consumers gain as they exhibit the same behaviour in all feasible contexts. The average welfare level becomes zero when the context goes beyond $\boldsymbol{x}_{46}$ because all types of consumers stop purchasing in the market.

2.3.3.4 The monopolist's optimal strategy, equilibrium contexts, and consumer welfare: Case 2

To examine whether the firm has an incentive to take advantage of consumers' ignorance by inducing supply distortions, the monopolist's profit-maximising responses to market demand and equilibrium contexts are studied in the case where products' prices and marginal costs are linearly associated with their attribute values. Consumer welfare under equilibrium contexts is also examined.

Consider the second case where $c_j = \delta_c \cdot s_j^m$ and $p_j = \delta_p \cdot s_j^m$ for all j, and the profit function is $\Pi(\boldsymbol{x}) = (\delta_p - \delta_c) \cdot \sum_j [s_j^m \cdot q_j^{D^*}(\boldsymbol{x})] - FC$. Suppose that fixed costs and δ_c are small enough that $\Pi(\boldsymbol{x}) \geq 0$ for all contexts ranging from $\boldsymbol{x}_1$ to $\boldsymbol{x}_{45}$. As consumers only choose one option after learning, the profit-maximising context in this case can be derived by comparing the attribute value of the *i*th option in the last context where it is chosen. For example, to obtain the benchmark profit-maximising context, one can simply compare $s_{x_3|\boldsymbol{x}_7}^m$, $s_{x_2|\boldsymbol{x}_{17}}^m$, and $s_{x_1|\boldsymbol{x}_{45}}^m$. The firm's optimal context choice in the benchmark model is $\boldsymbol{x}^* = \boldsymbol{x}_{45}$, the context just prior to the context that generates zero demand. In the proposed model, the firm confronted with consumers who experience zero or moderate degrees of social influence will also optimally produce at $\boldsymbol{x}^* = \boldsymbol{x}_{45}$. Interestingly, owing to deviation from optimality, the profit-maximising contexts in the market comprised completely of the (Φ = 1) type consumers are $\{\boldsymbol{x}_{35}, \boldsymbol{x}_{45}\}$. A graphical illustration that reveals the variation of profits is shown in Figure 5.

Apparently, in this case the firm is not incentivised to take advantage of consumers' bounded rationality by distorting contexts. More specifically, the firm in the proposed model does not have a monetary incentive to provide contexts with larger attribute values than the benchmark ones. Although the properties of the context, including the space between options and the numbers of options, matter, this estimated outcome occurs mainly because in the proposed model, the threshold criterion purely depends on the preferences$_U$, which limits the influence of false consensus on choice, and in turn leads to the same equilibrium context across models.

Moreover, according to Figure 4, all types of consumers gain the same level of welfare at $\boldsymbol{x_{45}}$ as they all choose x_1 in this context. This suggests that biased consumers are not worse off at equilibrium, even though in some circumstances they exhibit a compromise tendency and do not behave consistently. Therefore, although learning in the market per se is unable to totally eliminate the bias, it somewhat benefits false consensus consumers in the sense that it helps consumers choose the best options and avoid being exploited in the market equilibrium when the non-optimal option is strongly dominated.

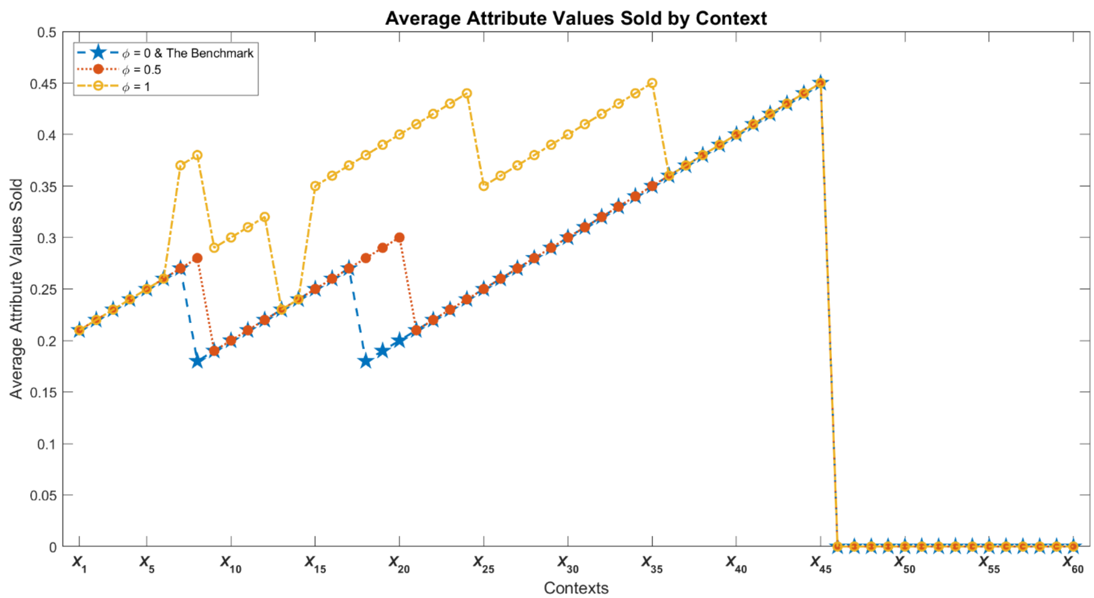

Figure 5. Attribute value of the demanded option by context. The deviation implies that consumers with Φ = 0.5 and 1 do not choose optimally under the corresponding context.

2.3.4 Simulation 3: Multi-agent scenario

Simulation 3 models a large number of heterogeneous consumer agents in an artificial market to explore the compromise effect phenomenon at an aggregate level. The simulation created an artificial market with fifty unit-demand consumer agents varying in the mean of their preferences$_{\mathrm{U}}$ distributions over the attribute *m*, i.e., the mean of

Beta(α_i^m, β_i^m). Fixing the variance of all PU_i^m at 0.0089[lxxix], the computer randomly generated values of the mean of each agent's preferences$_U$ from Beta (5, 5)[lxxx]. The computer redrew all values of the mean if there existed any mean that caused $PU_i^m(x)$ to be infinity or if there existed any $\alpha_i^m, \beta_i^m \leq 1$. In total, twenty sets of consumer preference$_U$ were sampled and the simulation randomly chose the 9th set for analysis. The set of fifty consumers' preferences$_U$ curves used in this subsection is shown in Figure 6.

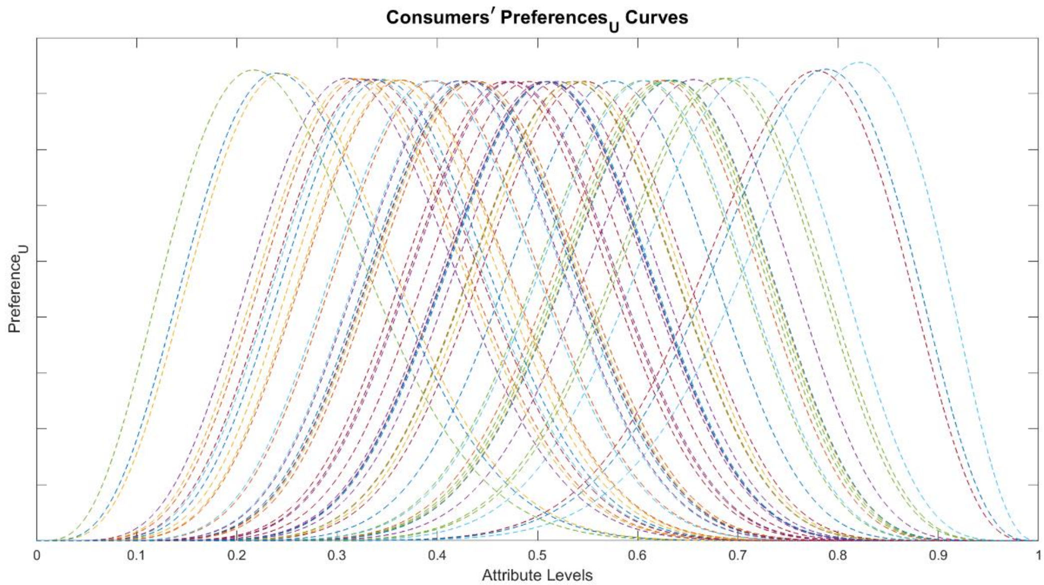

Figure 6. Consumers' preferences$_U$ curves in the simulated market. All distributions have variance of 0.0089 and different means randomly sampled from Beta (5, 5).

[lxxix] It is the variance of Beta (5, 15), the preference$_U$ distribution used in Simulation 1 and 2.

[lxxx] The authors note that in many real-world scenarios, the distribution of preference$_U$ in the population would be (positively) skewed. Examples include people's preferences for portion size, spiciness, cheese sharpness, etc. Yet, the overall finding and conclusion obtained in the case of skewed preference distribution largely resembles the one explored in Simulation 2. To avoid repetition, this subsection concentrates on investigating the market where consumer preference is symmetrically distributed.

The values of the rest of the parameters are identical to the ones in Simulation 2. False-consensus consumers have preferences$_{\text{N}}$ (PN^m) represented by a normal distribution, three possible values of the weighting parameter (Φ), which indicates the strength of the tendency to make market inferences during decision making, and a decision-making threshold (t) of 0.5. The computer simulation started from the context with $\boldsymbol{s}^{\boldsymbol{m}}_{\boldsymbol{x_1}}$ = {0.01, 0.11, 0.21, 0.31, 0.41} and ended at the context with $\boldsymbol{s}^{\boldsymbol{m}}_{\boldsymbol{x_{60}}}$ = {0.60, 0.70, 0.80, 0.90, 1.00}, with the attribute value of each option increasing by 0.01 units each time.

<u>2.3.4.1 The probability of opting out</u>

This subsection explores the probability of choosing nothing from the market offerings in both benchmark and the proposed models. Figure 7 shows the market quantity demanded for each option, including the outside option, under each feasible context. Note that consumers who do not choose from the market offerings are assumed to choose an outside option. Moreover, given that both benchmark models and the proposed models use the same value for the decision-making threshold, t, and the threshold criterion only involves value comparisons in preferences$_{\text{U}}$, the probability of opting out (i.e., choosing the outside option), is identical across both models, regardless of levels of Φ assumed in the proposed model.

In addition, examination of Figure 7 (a) reveals an approximately symmetric distribution of total market quantity demanded, with no consumer opting out in contexts between $\boldsymbol{x_{18}}$ and $\boldsymbol{x_{45}}$. This is because the mean of preferences$_{\text{U}}$ in the market follows a symmetric distribution, namely Beta (5, 5), centred at 0.5. Given that the difference in attribute values between adjacent products in the set, d, is fixed at 0.10 and n = 5, when the mean of $\boldsymbol{s}^{\boldsymbol{m}}_{\boldsymbol{x}}$ approaches 0.5 more consumers will choose to buy in the market since the threshold criterion is more likely to be met. By this token, if the mean of preferences$_{\text{U}}$ was drawn from a skewed distribution, the distribution of market quantity demanded would be skewed in the same direction. Since the threshold criterion is also employed in the benchmark model, total market quantities demanded in both models are identical under each context.

2.3.4.2 Quantity demanded estimated using different values of Φ and the compromise effect

This subsection discusses how estimated quantity demanded of a market alternative differs with values of Φ. In other words, this subsection is concerned with the effect of using marketing inferences during consumer choice. In addition, by comparing quantity demanded estimated by the proposed models (respectively with Φ = {0. 0.5, 1}) and the benchmark results, this subsection illustrates how the compromise effect arises and when it disappears.

The estimated quantity demanded of market offerings in the proposed model may be different from the benchmark results, although they exhibit a similar pattern. According to Figure 7(f), market quantity demanded of the fifth option, $q_5^{D^*}$, increases as the context moves from $\boldsymbol{x_1}$ to some points around $\boldsymbol{x_7}$[lxxxi], followed by a dramatical decline. Similarly, Figure 7(b) reveals that $q_1^{D^*}$ climbs gradually before the context reaches some points beyond $\boldsymbol{x_{53}}$, and then sharply drop. Note that x_5 gains more market demand than x_1 does until context is at around $\boldsymbol{x_{28}}$ and $\boldsymbol{x_{29}}$, and $q_1^{D^*}$outnumbers $q_5^{D^*}$ for all contexts beyond this point. These choice patterns appear in all models because, when the attribute values of market offerings are small overall, only a few consumers choose to buy and most of them are those who strictly prefer x_5 to x_1. As the context moves towards $\boldsymbol{x_{29}}$, more and more consumers opt in, while those who chose to buy when the context includes very small attribute values gradually switch to choosing other options over x_5 or opt out. Similarly, when the context approaches the high-end, more and more consumers will opt out and those who stay or start choosing to buy mostly are those who prefers x_1 over x_5.

The reason why $q^{D^*}(x_1|\boldsymbol{x})$ and $q^{D^*}(x_5|\boldsymbol{x})$ do not peak at $\boldsymbol{x_{28}}$ or $\boldsymbol{x_{29}}$ is that when the context is close to these two points, more and more consumers who choose to buy may prefer the middle option to the other options. Therefore, the curve of $q^{D^*}(x_3|\boldsymbol{x})$ seems to be symmetric, as shown in Figure 7(d). Note that these values do not exactly peak at $\boldsymbol{x_{30}}$

[lxxxi] Since the mean of each preferences$_U$ distribution is a random draw, the exact context that maximises $q_5^{D^*}$ depends on the results of the randomisation.

since mean of preferences$_U$ is a random variable which may not spread symmetrically in practice.

Moreover, due to the false consensus effect, the magnitude of $q^{D^*}(x_3|\boldsymbol{x})$ is higher in the proposed model than in the benchmark model when Φ = 0.5 and 1 (see Figure 7(d)). This implies the presence of a compromise effect in the proposed model under each of all contexts. More specifically, according to Figure 7 (d), the suboptimal, compromise option was more likely to be chosen as the average attribute values of market options approached 0.5 ($\boldsymbol{s}^{\boldsymbol{m}}_{\boldsymbol{x}_{29}}$ = {0.29, 0.39, 0.49, 0.59, 0.69}). This is because the bell-shaped distribution of the mean of preferences$_U$ causes the modes of preferences$_U$ to be distributed around 0.5. As a result, when the market context gets closer to $\boldsymbol{x}_{29}$, more and more consumers may choose to buy. Meanwhile, choice pairs that involve the middle option may be more likely to pass the threshold criterion, raising the possibility that false consensus consumers are affected by inaccurate estimation about their relative standings in the population during learning. As the market context moves away from $\boldsymbol{x}_{29}$, the compromise effect gets weaker.

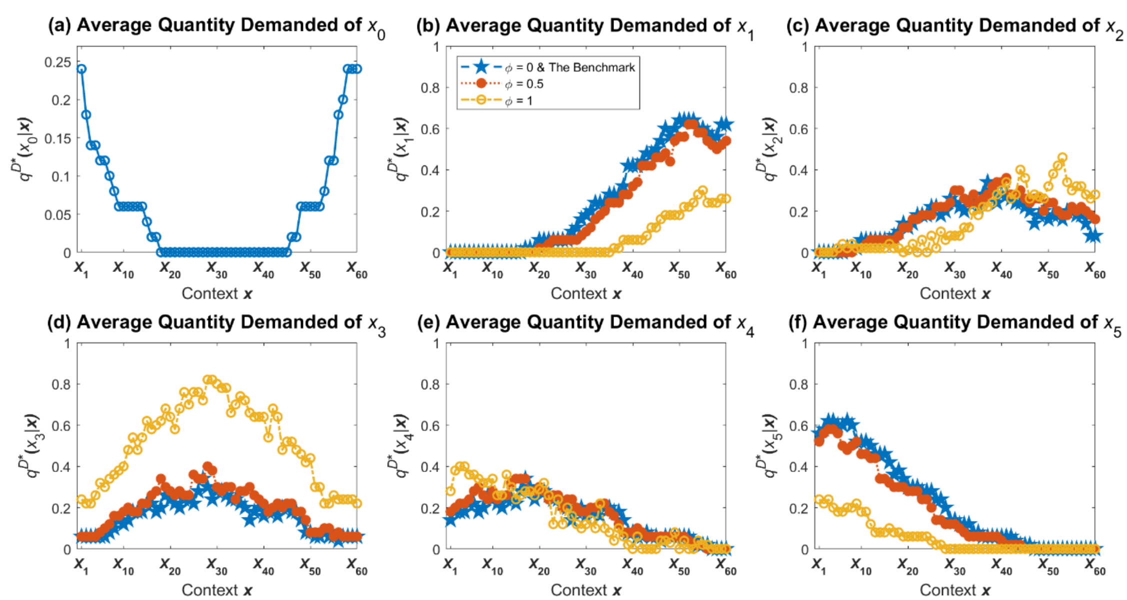

Figure 7. Market quantity demanded for each option by context, averaged across the number of consumers. The first figure indicates the proportion of consumers in the market who choose the outside option. The reverse of the graph implies the averaged market quantity demanded. Because of the constant threshold criterion, the curve is the same for all types of consumers.

2.3.4.3 Consumer welfare

This subsection explores average consumer welfare gained from consuming in the market as well as the effect of utility of an outside option, x_0, on average welfare. Consumer welfare under market equilibrium and a change in welfare due to distortion of market context will be addressed in the next two subsections, together with a discussion of the firm's optimal strategy.

Figure 8 (a) indicates that the mean level of welfare gained from choosing in the market is lowest under the first and the last context owing to an unsatisfactory choice set that causes a majority of consumers to opt out and choose nothing at all. Conversely, the highest mean welfare level is obtained under contexts around the middle of the space of feasible contexts. This may be because the attribute values of options there satisfy most consumers, even though a very high proportion of false consensus consumers may not choose the optimal option in the market due to bias. This also explains why the average welfare level gained in the extreme context can be effectively raised by increasing the utility of the outside option, as shown in Figure 8 (b), (c), and (d).

Furthermore, the deviation from optimality due to social influence is reflected in the gap in absolute welfare level between the benchmark results and the ones estimated by the proposed model with $\Phi = 0.5$ and 1. The greater the value of Φ, the larger the gap (see Figure 8 (a)). However, this welfare gap is unlikely to be reduced by increasing the utility of the outside option because the probability of opting out is the same for all types of consumers.

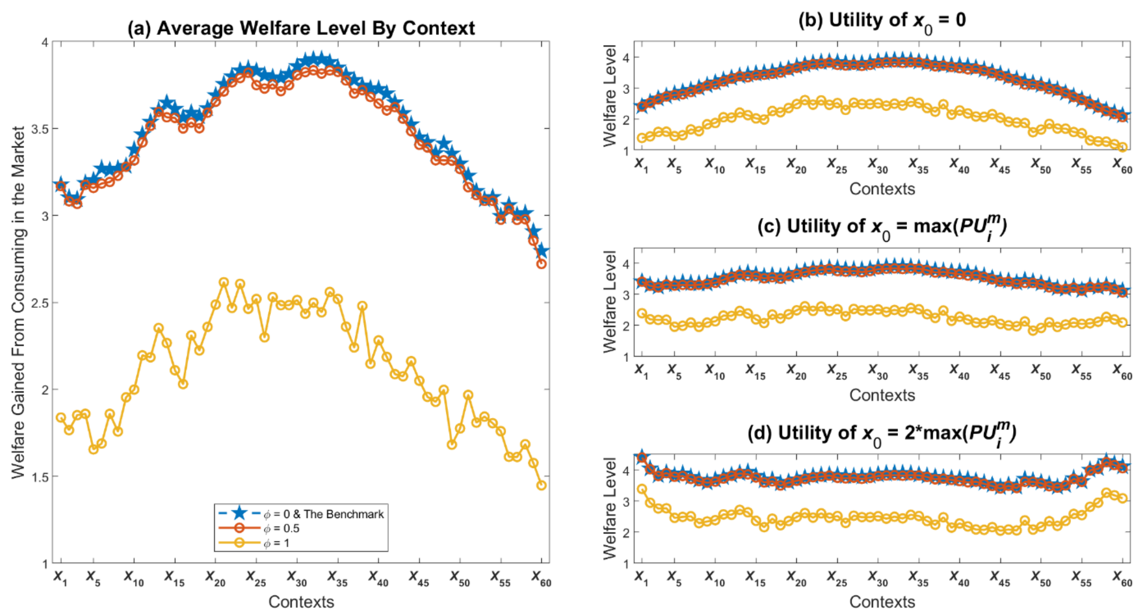

Figure 8. Consumer welfare under each context. (a) The average welfare level gained from choosing in the market. The context that maximises average welfare from consumption for the cases where Φ = {0, 0.5, 1} and the benchmark is {x_{33}, x_{32}, x_{21}, x_{33}}. In addition, the magnitude of the difference in welfare level between the proposed model and the benchmark results indicates the degree of deviation from optimality.; (b), (c), and (d) Average consumer welfare with three possible levels of utility of the outside option. The difference among these three figures implies that the high probability of choosing outside may improve welfare level considerably if utility of the outside option raises to a certain level. Therefore, the shapes of all curves reverses as $PU_i(x_0)$ increases.

<u>2.3.4.4 The monopolist's optimal strategy, equilibrium contexts, and consumer welfare: Case 1</u>

The monopolist's best response to market demands and equilibrium contexts in the case where prices and marginal costs are constant regardless of products' attribute values is now investigated. Consumer welfare market equilibrium context in this case is also explored.

Theoretically speaking, if selling a product with more attribute value cannot increase profits, the firm's profit-maximising strategy of context setting will be to sell as many products as possible, instead of trying to promote a particular product at expense of losing demand for other offerings. Therefore, the profit-maximising contexts in the case of a fixed *p* and *c* should be the contexts that minimise consumer' probability of opting out. Consistent with information shown in Figure 7 (a), the monopolist is best-off producing at $\boldsymbol{x}^*$ = {$\boldsymbol{x}_{18}$, $\boldsymbol{x}_{19}$, …, $\boldsymbol{x}_{45}$} in the benchmark and proposed models as long as the profit is positive. The reason why the set of profit-maximising contexts does not differ with models (i.e., the benchmark versus the proposed model) and the values of Φ used in estimation is because the assumed value of the decision threshold is identical for both models and the value of Φ cannot affect results of threshold assessment in pairwise comparison. Therefore, in line with theoretical predictions and previous findings, when profits do not depend on which option is chosen, the firm has no incentive to take advantage of consumer bias.

In addition, the set of market equilibrium contexts $\boldsymbol{x}^*$ contains a socially optimal outcome in both models. As indicated in Figure 8 (a), mean consumer welfare in the benchmark model and the proposed model with Φ = 0, 0.5, and 1 is highest when context is {$\boldsymbol{x}_{33}$, $\boldsymbol{x}_{32}$, $\boldsymbol{x}_{21}$, $\boldsymbol{x}_{33}$}, respectively. This result demonstrates that a strong compromise effect may exist in equilibrium even when complete learning is possible and when consumers choose rationally based on their learned preferences.

2.3.4.5 The monopolist's optimal strategy, equilibrium contexts, and consumer welfare: Case 2

Now consider the second case, where price and marginal cost are linearly and positively associated with values of attribute *m* of products, $\boldsymbol{s}_{\boldsymbol{x}}^{\boldsymbol{m}}$. As before, the firm's best responses to market demand and equilibrium contexts in this new case are investigated in this subsection. In addition, if manipulation of market context is likely to happen, a change in consumer welfare due to manipulation is examined for different types (with respect to Φ) of consumers.

Recall that the profit function in this case is $\Pi(\boldsymbol{x}) = \left(\delta_p - \delta_c\right) \cdot \sum_j \left[s_j^m \cdot q_j^{D^*}(\boldsymbol{x})\right] - FC$. Accordingly, for a monopolist to earn positive profits, the term $\left(\delta_p - \delta_c\right)$ should be positive. Suppose that the fixed costs are small enough that $\Pi(\boldsymbol{x}) \geq 0$ for all contexts that generate positive demand. Profit-maximising contexts are $\boldsymbol{x}^*$ = {$\boldsymbol{x}_{45}$, $\boldsymbol{x}_{47}$, $\boldsymbol{x}_{47}$} respectively for Φ = {0, 0.5, 1} in the proposed model and the benchmark $\boldsymbol{x}^*$ = $\boldsymbol{x}_{45}$. This result is in line with the theoretical predictions, which suggest that the optimal context may be closer to the upper bound of the set of all feasible contexts than the one that maximises total quantity demanded (i.e., $\boldsymbol{x}^*$ in the previous case) will be when prices and marginal costs are linearly dependent on attribute values of products. More specifically, when $\Phi = 0.5$ and 1, an increase in the average attribute value sold compensates for the very small loss in overall quantity demanded. This strongly incentivises the firm to raise the overall attribute values of market offerings in the proposed model. Hence, it is predicted that in this case, there is a strong monetary incentive for firms to distort the informational content of product lines to take advantage of consumer bias.

However, the welfare impact of context distortion is negative for all types of consumers. As shown in Figure 8 (a), for consumers with $\Phi = 0.5$, the average consumer welfare gained from consuming in the market declines from 3.4061 at $\boldsymbol{x}_{45}$ to 3.3169 at $\boldsymbol{x}_{47}$. Similarly, the average welfare for the ($\Phi = 1$) type consumers decreases from 2.0493 at $\boldsymbol{x}_{45}$ to 1.9282 at $\boldsymbol{x}_{47}$. This suggests that profit-induced context distortion may psychologically harm biased consumers in a market equilibrium, despite consumers' efforts in learning their preferences and choosing the optimal products.

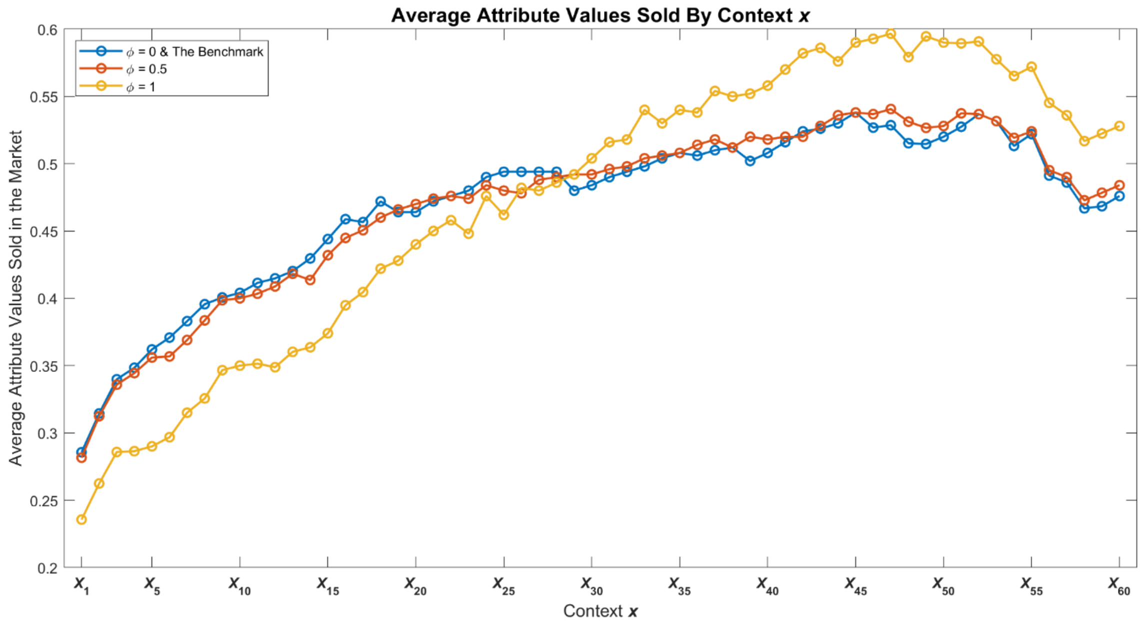

Figure 9. Attribute values sold by context, averaged over fifty sampled consumers. Comparing across the three lines, the sizes of profits are similar when consumers suffer zero to moderate degree of social influences. This reflects the similar simulated choices made by the benchmark consumers and by consumers with Φ = 0 and 0.5. For the market constituted by consumers with Φ = 1, the profits earned under the first half of contexts, namely contexts prior to $\boldsymbol{x}_{\mathbf{28}}$, are strictly lower than benchmark profit levels. This is because under those contexts, biased consumers tend to choose x_3 and x_4 rather than x_5, which is the option that generates more profits. The situation reverses under the second half of the contexts (i.e., those with larger values) because biased consumers are less likely to choose x_1.

2.4 Summary

This chapter proposes a model of context effects to explain and predict consumer behaviour with preference learning. The model incorporates RRT, the rank-order decision rule, and the false-consensus effect. Most of the assumptions held in the RRT are retained in the proposed model, including the assumption that final purchasing choices are driven not directly by the pre-existing preference$_{\mathrm{U}}$, but by the context-specific preference$_{\mathrm{I}}$, which is formed by outcomes of preference learning in the market. More concretely, the model inherits the assumption that the conscious inaccessibility of values of the preference$_{\mathrm{U}}$ may motivate decision makers to learn their preference through a series of pairwise comparisons between market offerings. To prevent people from choosing extremely undesirable products, it is also assumed that, for each pair, a choice will be made only if the difference in the preference$_U$ between alternatives exceeds a certain threshold value. If the pair passes the threshold criterion, it is uniquely assumed by the proposed model that choice probability of a paired alternative is co-governed by internal signals (preference$_U$) and the market inference (the rank-order decision rule). Since false consensus individuals tend to generate flawed estimates of their relative standing in the population's taste distribution, they may be biased to believe that the middle option fits them the most during preference learning. Consequently, the presence of this type of consumer potentially contributes to the observations of the compromise behaviour in the market. In other words, according to the proposed model, the misperception of one's relative position in the population distribution may incur errors in preference learning as consumers adopt the rank-order decision rule, and this is the cause of compromise behaviour.

However, the involvement of errors in making market inferences does not necessarily mean consumers are irrational. Believing social norms and product lines are informative and thus attend to them can arguably be viewed as a rational move. Therefore, the proposed model indeed suggests that rationally using a strategy to assist decision making can lead to seemingly irrational behaviour.

The proposed model is complex. However, each aspect of it is essential and thus cannot be omitted. Firstly, as with RRT, the distinction between preferences$_U$ and preferences$_I$ is crucial in modelling human behaviour, whereby preferences$_U$ are important in explaining natural individual differences whereas preferences$_I$ reflect environmental influences. Therefore, it is necessary to include these two preferences in the model. More concretely, having a preference$_U$ means that there must be two parameters that jointly define preferences$_U$'s properties (as a Beta distribution). Moreover, having preferences$_I$ requires us to understand the learning process because decision makers need to form these preferences$_I$.

Incorporating the idea from psychology literature that a preference$_U$ can only underpin ordinal information about it, the learning of preferences through a series of pairwise comparisons between choice alternatives is a necessary assumption. Otherwise, preferences$_U$ cannot exert their influences on choice. Moreover, the threshold parameter that appears in the first stage of learning is also necessary, because it is a reasonable and simple mechanism that prevents people from choosing extremely undesirable products.

In addition, the involvement of social influences in the pairwise learning process is assumed based on extensive behavioural evidence that shows the effect of social norms on decision making (e.g., Herman, Polivy, Pliner, & Vartanian, 2015). Without it, the model may be less realistic. In addition, the reason why the proposed model adopts the rank-order decision rule to embody the idea of social influences is because it has been proven effective in explaining the compromise effect as well as decisions under preference uncertainty (see Section 1.3.3 and 1.3.4 for more information). Yet, as the rank-order decision rule does not explicitly specify how decision makers estimate their relative standings in the population's taste distribution, the notion of the false consensus effect must be incorporated into the proposed model (for review, see Section 1.3.4). To account for individual differences in resisting social influences, the parameter Φ that controls the tendency to use the rank-order decision rule is necessary.

Last, to reduce complexity of the proposed model, the joint effect of preferences$_U$ and preferences$_N$ is computed by the convex combination approach, rather than any other more complex integration. Similarly, the decision made in each pair is computed based on a simple but useful softmax function. Finally, the way that the preferences$_I$ are formed from outcomes of preference learning in the proposed model is directly inherited from RRT. It is one of the most essential parts of the model as the existence of preferences$_I$ is core of the model. Altogether, based on the above-mentioned reasons, it is believed that each part of the proposed model is essential. Removing any of them will decrease the descriptive power of it.

The predictions of the proposed model were derived via computer simulations. In the single agent case, the results suggest that decision makers who do not use the rank-order decision rule at all, i.e., whose Φ = 0, behave in the same manner as the benchmark consumer[lxxxii] does. This implies that unbiased preference learning may lead to optimal outcome, even in presence of uncertainty about the magnitudes of the preference$_U$. However, consumers with Φ = 0.5 and 1 may exhibit a strong tendency to choose the options close to the middle, with consumers who use the biased market inference more frequently (i.e., Φ = 1 versus 0.5) showing more biases towards the middle option. In addition, choice reversals are observed for these two types of consumers in some contexts. This may potentially prompt an outside observer to conclude that the preference$_U$ is context-dependent, although, in effect, it is the relative ranks of available products that affects relative attractiveness without changing preferences$_U$. Also note that choice reversal and the propensity to choose the compromise option only appear in some contexts. For a context that is very unfavourable, even consumers with Φ = 1 would act consistently with the benchmark agent. Altogether, the results on the demand side demonstrates that when the value of the preference$_U$ of the middle option is high enough, the perceived locations in the population's distribution play a crucial role in decision making by biasing outcomes of preference learning. Yet, if the middle option is too unattractive in terms of its absolute attribute value, the existence of the context-independent preference$_U$ and market experience are sufficient to prevent biased consumers from endlessly exhibiting the compromise behaviour.

With regard to the firm's profit-maximising strategy, when prices and marginal costs are constant across all feasible products, it is found that the monopolist has no incentive to manipulate the market context. Likewise, in the case where attribute values of sold products are direct inputs to the profit function, i.e., the firm can benefit from selling products with larger attribute values, the simulation results still do not find evidence of incentives for context distortion. This outcome can be attributed to the fact that in the proposed model, the threshold criterion purely depends on the preferences$_U$, which limits social influence on choice, and in turn leads to the same equilibrium context

---

[lxxxii] The benchmark model is not a special case of the proposed model. It assumes consumers know their own preferences$_U$ values, and thus can make a purchase decision accordingly.

across models. Moreover, given that choice in the equilibrium context is identical across all consumer types, biased consumers are not found to be psychologically worse off. Therefore, although learning in the market *per se* cannot beat the bias, it may benefit false-consensus consumers by helping them choose the better options and avoid being exploited in the market equilibrium.

The present chapter also conducted a simulation of a multiple-agent scenario. The overall results obtained from fifty randomly sampled, unit-demand consumers are roughly consistent with the theoretical predictions. Firstly, as found in the previous case, consumers who do not use market information at all, i.e., whose $\Phi = 0$, make the same choice as the benchmark consumer does in all feasible contexts. In contrast, consumers with $\Phi = 0.5$ and 1, on average are more likely to choose the middle option than the benchmark consumer does, with the degree of the bias increasing with the value of Φ. In addition, for these two types of consumer, the compromise tendency gets weaker as the market context moves towards the boundary of the space of feasible contexts. This is possibly because more consumers choose to opt out when context becomes more extreme, which shrinks the effect of the preference$_N$ on choice since the rank-order decision rule is only used at the second stage of the pairwise comparison.

The results on the producer side are largely consistent with the theoretical predictions. Given that the likelihood of opting out is the same for all types of consumer, the set of the profit-maximising menus is the same when prices and marginal costs are constant. Since the mean consumer welfare gained in the market differs across types of consumers, the socially optimal equilibrium context may not be identical. Additionally, even at the socially optimal equilibrium point, consumers who suffer the strongest social influences, i.e., $\Phi = 1$, gain lowest welfare due to their biased behaviour.

In another case, where price and marginal cost are positively associated with the attribute values of products, the results suggest that a monopolist may be monetarily incentivised to offer products with higher attribute values to biased consumers, i.e., those with $\Phi = 0.5$ and 1, than to the benchmark agent. Yet, this finding is highly dependent on the preference$_U$ distributions sampled, as the results in the single-agent case clearly illustrate the possibility that distortion of informational content of product lines may not occur. Moreover, the welfare impact of context distortion is negative for consumers with

Φ = 0.5 and 1, suggesting that both consumers experience a welfare loss resulting from being offered a less favourable context.

In a nutshell, the results demonstrate that the compromise effect may exist in market equilibrium, even when preference learning is possible and consumers choose rationally based on their learned preferences. It is also found that if the market comprises many heterogenous, false-consensus agents, it can be profitable for a monopolist to distort menus, but this conclusion strongly depends on the settings of parameter values and market contexts. Combining these two points, the simulation results suggest that the existence of biased consumers may be the reason why compromise behaviour is often observed in the market and why the possibility of context distortion should be of great concern. Fortunately, the results also find that there is a limit to menu manipulation since unpleasant signals sent from the preference$_U$ will become more salient as market offerings get more unsatisfactory. In addition, context distortion does not always harm consumers as they may choose less optimally in the benchmark equilibrium context. However, based on these results, it seems that the influence of biased estimation about social norms is not strong enough to exert a powerful impact on the downstream behaviour. At least, it is observed that a very high value of Φ, e.g., Φ = 1, is required to show the effect. This limitation motivates development of an alternative model which allows social norms to have a more direct and independent effect on the final purchasing choice. The details of the second model and its predictions will be presented in the next chapter.

Chapter 3: Model 2

This chapter develops an alternative model of context effects in consumer choice under repeated-purchase conditions. As in Chapter 2, the purpose is to model both (a) the role of market context in consumers' judgement and decision making and (b) a firm's best response to (biased) consumer behaviour. To recap, the key research questions that the thesis aim to address are: (1) If consumers are endowed with a context-independent preference$_U$, why do they behave as if they are biased towards the compromise options? (2) Under what conditions will the compromise effect persist or be eliminated in market equilibrium when consumers are experienced? Also, is there any other factor that contributes to the compromise effect? (3) If increasing the attribute values of products is profitable, what is a firm's optimal strategy in setting the product line, and how does this strategic reaction in turn affect consumers' welfare in equilibrium?

As discussed in Chapter 2, there are some limitations to the first model that make its behaviour slightly different from initial expectations. The main problem of the first model is that the decision about whether to buy (at all) in the market solely depends on innate, context-independent preferences$_U$. As attribute values of available products get larger (i.e., the choice context moves towards the upper extreme of the attribute space), the decision maker becomes more likely to choose nothing at all in pairwise comparisons. This leads consumers to eventually choose the first option from the ordered choice set with certainty (i.e., with a probability of 100%) when market context goes beyond a critical point[lxxxiii]. Thus people's market-based inferences play a very limited role in final purchase choice when there are too few available options satisfying the preference$_U$. This in turn gives the firm no incentive to distort the market context to take advantage of biased consumers. Another problem of the first model lies in its assumption that consumers will

[lxxxiii] When the preference$_U$ of the first option is the only one that passes the threshold, the first option will be the most frequently chosen one during learning and therefore obtains the highest preference$_I$, compared to other options. Social norms do not have any effect on the final decision in this case.

never adjust their purchase decision once they have formed their own inferred preferences. This means that they will always choose the same product for all remaining periods, which is inconsistent with consumer behaviour as documented in the literature[lxxxiv]. Finally, the assumption of pairwise comparisons may not hold in reality. In particular, for products varying in portion size, it is extremely unlikely that consumers will in reality compare two products simultaneously. This fact limits the applicability and generalisability of the first model.

These limitations motivated development of a simpler one-dimensional model that abandons the strong assumption of pairwise comparisons, while retaining the assumptions about preferences$_U$ and norm-based preferences with false consensus. More specifically, the new model employs a classical explore-exploit framework and assumes that false-consensus consumers will choose either on the basis of preferences$_N$ (i.e., exploration), or on the basis of their inferred preference, which is formed by their past choices (i.e., exploitation).

A key new aspect of the new model is that choice will lead to subsequent *inhibition* of both chosen and (some) non-chosen options. This version of “inhibition of return” is a psychological mechanism that prevents people from immediately re-choosing products that are known to be unfavourable until enough time has passed for the inhibition to decay. Regardless of which strategy (explore vs exploit) is adopted, on any given occasion consumers will choose only from options that are inhibited less than some threshold values. These threshold values may in turn be influenced by the preferences$_N$, such that inhibited items that are believed to be desirable (according to market inferences) will re-enter an individual’s choice set sooner than will equally-inhibited items that are not believed to be desirable. Note that preferences$_N$ can therefore affect choice in two quite

---

[lxxxiv] It is suggested in the consumer research literature that consumers tend to choose variety with the objective of counteracting physiological and psychological satiation from repeated choices (Kahn, Kalwani, & Morrison, 1986; Levav & Zhu, 2009; Sevilla, Lu, & Kahn, 2019).

different ways — by determining choice probability in exploration trials, but also by influencing how quickly inhibited items become re-eligible for choice.

More specifically, after every choice that is made in the market, inhibition is added to that option and to all options that are inevitably less favourable than it. The amount of this inhibition of return is determined by the preference$_U$ value of the chosen option and by two other exogenously determined parameters. The amount of inhibition, and also the memory of choice history, are assumed to decay over time with a constant rate. Overall, by allowing social norms to influence decision thresholds, choice is made based on preferences$_U$ and the social norm more directly. Moreover, in this model, consumers never stop learning and updating their inferred preferences, and therefore their choices may change over time.

A concrete example of the proposed consumer decision making process is provided in the next section for illustrative purposes, followed by a review of the literature relating to two main additional assumptions of this alternative model – the explore-exploit trade-off, and inhibition of return. Then, a formal description of the new model as well as the firm's profit-maximisation problem will be presented. The chapter will close by presenting a series of computer simulations, along with a discussion of estimated consumer behaviour and market equilibrium.

3.1 An Example

Consider two representative consumers, Alex and Sam, who are each endowed with a preference$_U$ for sizes of a soft drink but cannot consciously access the relevant absolute magnitudes. One day they, separately, go to a newly-opened shop in which soft drinks are labelled in terms of their sizes on an ordinal scale (e.g., small, medium, and large) and displayed in order. Alex and Sam do not have any prior experience of purchasing or consuming soft drinks sold in that shop.

Suppose that Alex always chooses based on his/her feelings, i.e., preferences$_U$, and never makes market-based inferences about social norms. At the first time of visiting, Alex will randomly choose one from the set of available soft drinks ("exploration"). After consuming the chosen drink, Alex will have a feeling about its favourability (based on a preference$_U$) and based on that feeling will unconsciously assign inhibition to it. Items with lower preference$_U$ values will receive more inhibition. The magnitude of this inhibition of return represents the degree that he/she is unwilling to choose it again. Moreover, if Alex is able to identify whether the size of the choice is too large or small for him/her, the inhibition will be spread to options that are perceived as more unfavourable (i.e., all drinks that are larger or smaller than the choice). However, if the choice is good enough[lxxxv], Alex may fail to consider other drinks' desirability and hence inhibition of return will not spread out, while the chosen one is nonetheless inhibited.

At the second time of visiting, Alex may either choose randomly again (exploration) or choose based on the preference inferred from his/her past choices (exploitation). Note that no matter which strategy is adopted, products will be considered only when their inhibition is less than a threshold value. Because Alex is not sensitive to social norms, these threshold values will be constant for all items. If all soft drinks in the shop are inhibited above threshold levels, Alex will make no purchase. The post-choice stage will be same as in the first time period, unless Alex decides not to buy in the shop at all[lxxxvi]. Further, at the beginning of each shopping period, memory of inhibition of return and choice history decay by a constant amount. Following this proposed decision-making process, it is predicted that if all soft drinks in the shop are very undesirable with respect to the preference$_U$, Alex will nonetheless end up choosing the least unfavourable one occasionally, with other products even more rarely being chosen.

Consider now another consumer, Sam. Suppose that Sam tends to adopt a matching strategy when deciding whether to purchase a soft drink. The use of the matching strategy indicates that Sam believes that the product line is designed to meet the whole population's tastes on the soft drinks, and therefore that the relative rank of the most desirable (for Sam) product in the product line should match his/her rank in the population's taste distribution. In other words, if Sam believes himself/herself to fall at

[lxxxv] That is, Alex cannot consciously recognise whether he/she prefers a larger or smaller option to the choice as it is so satisfactory in terms of the preference$_U$. Importantly, the "good enough" choice is not necessarily the one that maximises satisfaction of the preference$_U$.

[lxxxvi] In this case, none of the options will be inhibited in that round since no choice is made.

the *n*th percentile of the population distribution of tastes, Sam will believe that the product at the *n*th percentile of the distribution of market options will be the product that is most suitable for himself/herself. Yet, the model further assumes that Sam suffers a certain degree of false consensus effect when evaluating his/her own position, such that he/she is prone to estimating that his/her preference$_U$ is located at around the 50th percentile of the population. Accordingly, Sam may think that he/she most prefers the option with a relative rank of 50% and that he/she will prefer extreme options least. These beliefs will influence Sam's exploration-based choices.

At the first time of visiting the shop, Sam will choose in proportion to his/her beliefs about the social norms (i.e., choices made by others), which indicate that the medium size is the most favourable one and the extreme sizes of soft drinks are the least. This makes Sam most likely to choose the median option at first (in the first exploration-based choice). Similarly to Alex, in the post-choice stage Sam will unconsciously assign an inhibition to the chosen drink based on his/her preference$_U$. Again, whether the inhibition of return will be spread to other options depends on Sam's ability to identify the less favourable ones. If the choice is satisfactory, Sam may be unable to recognise whether the chosen drink is too large or too small in its size and therefore cannot make a sensible inference about other products' desirability.

On the second visit, Sam will either explore again, in which case choice probability will be based on the perceived social norms, or exploit based on past choices. Like Alex, Sam will only consider soft drinks with inhibition values less than threshold values. In other words, Sam will not choose options that he/she has very negative impressions of. However, due to the tendency to use the matching strategy, the threshold value is assumed to be increased, to various extents, by the false consensus effect. This is because to false consensus consumers, middle options are optimal products as suggested by market inferences. More specifically, Sam will be more tolerant of the middle-size drink, in the sense that he/she may still consider the middle option even when he/she has a negative impression on it. Sam is still more tolerant of extreme options than Alex, even though the extreme options are less likely to be considered than the medium one, due to their lower threshold values (i.e., inhibition must be smaller for extreme items to be considered for choice). This is because Sam tends to believe that the product line is reasonably designed to capture preferences – Sam's assumption is that none of the soft drink sizes presented in the shop are definitely deficient. Consequently, Sam is not only more likely to exhibit compromise behaviour, but also more likely to make purchases at all.

If observers infer Alex and Sam's preference$_U$ simply based on their choices under various market contexts, they may easily conclude that Sam behaves less rationally than Alex since Sam shows a great tendency to choose the middle option, regardless of its attribute value. However, this claim ignores the fact that both Alex and Sam cannot access the absolute magnitude of their preferences$_U$ and Sam, in fact, is the one that deliberately attempts to adopt a (rational) strategy in order to make a better decision. A diagram that summarises the proposed model is shown below (Figure 10).

Note that false consensus effects and social norm effects are intertwined in models such as the present one. Intuitively, a consumer might have a tendency to choose compromise options for reasons that are psychologically distinct but which are not distinguished in the present formulation of the model. Thus a consumer might believe that other people tend to choose extreme options less frequently than middle options (unimodal social norm), and also accurately estimate his/her own relative ranked position within the population. Alternatively, the consumer might believe that population choices are uniformly distributed over market options, but believe that his/her own relative ranked position within the population is closer to the median than it truly is. Either of these possibilities could give rise to compromise effects; the current implementation is neutral between them.

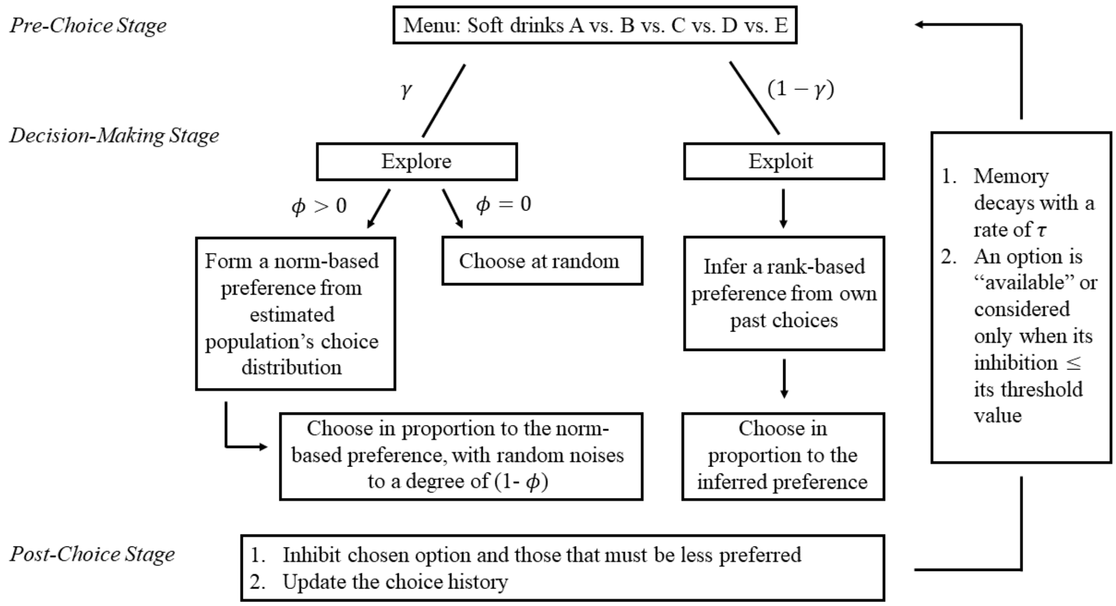

Figure 10. A diagram of the second model. The parameter γ refers to the likelihood of exploring, ϕ denotes the tendency to use the matching strategy, and τ indicates the rate of memory decay. In the example, Alex has ϕ = 0 and Sam's ϕ lies in an interval, $1 \geq \phi > 0$.

3.2 Assumptions

The novel features of this second model (as compared with the model presented in the previous chapter) are the explore-exploit framework and inhibition of return. Both concepts have been explored extensively and received substantial empirical support in literature. This subsection will review literature relevant to these two assumptions and explain how the assumptions were derived.

Searching is a commonly observed phenomenon in both the human and animal world and has been studied in fields such as behavioural science (Stojić, Schulz, Analytis, & Speekenbrink, 2020; Wu, Schulz, Speekenbrink, Nelson, & Meder, 2018), and biology and neuroscience (Blanco, Love, Cooper, McGeary, Knopik, & Maddox, 2015; Cohen, McClure, & Yu, 2007; Mobbs et al., 2013). While trying untested options or looking for new opportunities may bring uncertain outcomes, gathering information about the environment is of vital importance in terms of long-term reward maximisation (Schulz &

Gershman, 2019). The trade-off between exploiting familiar, relatively riskless options and seeking novel, more risky options frames the explore-exploit dilemma.

Existing research on the explore-exploit dilemma mainly focuses on exploratory behaviour, leaving the algorithms of exploitation relatively less well specified in the literature. Normally, exploitation refers to the choice of options with (possibly highest) known rewards (Feng, Wang, Zarnescu, & Wilson, 2021). Integrating with RRT, the proposed model extends the idea of exploitation to assume that consumers exploit by choosing based on their inferred preferences, which are formed based on the relative frequency with which an option has been chosen in the past. Note that since it is also assumed that people are insensitive to absolute attribute values, the inferred preference is completely rank-based. It is therefore an option's relative rank being exploited, not its attribute value or any other features.

In contrast to exploitation, decision-making in exploration trials has been widely studied. Exploration strategies can generally be classified into three groups[lxxxvii], of which the first two belong to the family of uncertainty-guided exploration (Wilson, Geana, White, Ludvig, & Cohen, 2014). The first group is called random exploration, which describes exploration as an injection of randomness into choice that enables choice stochasticity to reflect one's uncertainty about the outcomes (Schulz & Gershman, 2019). For example, options can be sampled randomly (Schulz et al., 2019) or in proportion to their respective (expected) values (Schulz & Gershman, 2019). The second group is termed directed exploration, where exploration is guided by assigning an uncertainty bonus to each available option or action (Auer, 2002). The inflated (by uncertainty) expected values channel sampling into more uncertain or unfamiliar options (Kakade &

[lxxxvii] There is a separate line of research discussing the rate of exploration. For example, Riefer, Prior, Blair, Pavey, and Love (2017) empirically examine consumer explore-exploit trade-offs over time and report that the longer consumers exploited a product, the less likely they were to explore, possibly due to coherence maximisation. Yet, in the second model, the explore-exploit parameter is a free parameter, which means that there is no specific constraint on it. Therefore, the current review of research on human exploration focuses on strategy rather than timing.

Dayan, 2002). Consistent with the theory, experimental findings (e.g., Gershman, 2018; Schulz, Wu, Ruggeri, & Meder, 2019; Wilson, Geana, White, Ludvig, & Cohen, 2014) confirms that human may adopt both types of uncertainty-guided heuristic.

The third, and more recently proposed, group is belief-guided exploration, which states that searching relies on structured knowledge or beliefs about the environment (Dayan & Berridge, 2014; Markant & Gureckis, 2014; Otto, Knox, Markman & Love, 2014). More specifically, it suggests that people leverage information about the structure of the environment to decide where and when to explore. In line with the theory, Schulz, Bhui, Love, Brier, Todd, and Gershman (2019) find that consumers buy from novel restaurants by adaptively referring to signals of restaurant quality and may generalise their experiences across the same type of restaurants. This finding is supported by Stojić, Schulz, Analytis, and Speekenbrink (2020), who show experimentally that knowledge about the outcome that an option's features are associated with generalise to other unexplored options. Therefore, they argue that whether human exhibit novelty seeking (or avoidance) depends on whether the novel option's features resemble those of highly rewarding (or nonrewarding) options. Furthermore, both Schulz, Bhui, Love, Brier, Todd, and Gershman (2019) and Stojić, Schulz, Analytis, and Speekenbrink (2020) observe some degree of directed exploration from people. This indicates that these exploration strategies are not mutually exclusive – decision makers may simultaneously consider prospective rewards and the degree of uncertainty.

Extending the existing literature, the belief that menus are designed to capture the population's taste distribution may predispose consumers to relate a product's relative rank to its subjective desirability. As a result, they are expected to be more inclined to adopt belief-guided exploration, which, by assumption, involves choosing proportionally to norm-based preferences with some degree of random error. Moreover, in the proposed model, menus may be local and small, and therefore exploration does not necessarily mean people try out an option that is never chosen. Instead, exploration here refers to the situation when choice is not purely based on past choice. Hence, in explore trials, consumers attempt to choose options that may surprise them, even though they may have tasted it.

However, modelling exploitative and exploratory behaviour in this way makes preferences$_U$, and thus utility gained from consumption, ineffective in decision making, potentially leading consumers to repeatedly choose inferior options in both trials. Information gathered from experience or feedback about one's previous behaviour is frequently used in human and animal search. Taking the example of food foraging: whether to return to a place that has been examined greatly depends on whether the locus of a food source in that area is found to be exhausted (Klein, 2000). In the visual search literature, a mechanism that encourages orienting towards a novel or significant stimulus by biasing attention and/or eye movements from an (overly) exposed stimulus is called inhibition of return (Klein & MacInnes, 1999; Posner, Rafal, Choate, & Vaughan, 1985; Tipper, Weaver, Jerreat, & Burak, 1994).

The term "inhibition of return" comes from an experimental finding of Posner and Cohen (1984), who showed that participants responded more quickly to targets presenting at previously cued locations than to targets at uncued locations as long as the time between the cue and the target, i.e., stimulus onset asynchrony, was short. This phenomenon, however, reversed when stimulus onset asynchrony increased, as evidenced by longer response times for cued targets than for uncued targets (Posner & Cohen, 1984). This finding suggests that the early facilitation, possibly thanks to the rapid and correct stimulus information at cued locations, is followed by a long-lasting inhibitory effect reflecting the removal of attention from that location (Bennett & Pratt, 2001; Klein, 2000; Satel, Wang, Trappenberg, & Klein, 2011; Taylor & Klein, 1998). In short, inhibition of return describes the phenomenon whereby attention is slow to return to an originally attended location.

The implication of inhibition of return is that people tend to restrain themselves from revisiting previously examined places, especially those that do not produce valuable outcomes for a long time. Following this insight, the second model assumes that consumers may label options they have chosen and be disinclined to reconsider them. The size of the inhibitory effect in the model is computed as an inverse function of one's own preference$_U$ and is assumed to accumulate over time. It is therefore expected that the worse the consumption experience is, the more likely the chosen option is to be inhibited, regardless of its relative rank in the set. Furthermore, inhibition of return is found to spread to adjacent locations in a graded manner (Birmingham, Visser, Snyder, & Kingstone, 2007; Christie, Hilchey, & Klein, 2013; Taylor, Chan, Bennett, & Pratt, 2015). Exploiting this observation, the proposed model assumes that options that are conjectured to be worse than the chosen one will suffer inhibition of return following the choice.

3.3 Theory

3.3.1 Model description

A formal description of the model is presented below[lxxxviii]. Let $\boldsymbol{x} = \{x_1, x_2, \ldots, x_n\}$ denote a finite set of *n* market alternatives ordered in terms of magnitudes of a common, quantifiable attribute *m*, and let $s_j^m \in [0,1]$ denote a normalised attribute value of an option x_j on an attribute *m*. A unit-demand consumer *i*'s preference$_U$ of absolute values of an attribute *m*, PU_i^m, is represented by the probability density function of a beta distribution, $\text{Beta}(\alpha_i^m, \beta_i^m)$, with $\alpha_i^m, \beta_i^m > 1$ to guarantee unimodality. Note that the exact value of PU_i is consciously inaccessible and consumers do not have any *ex ante* knowledge or experience of available products $\boldsymbol{x}$ and the technological feasible set *X*.

Moreover, as noted before, the norm-based preference, $PN^m(x|\boldsymbol{x})$, refers to a (context-dependent) preference relation acquired from the product line $\boldsymbol{x}$ and one's perceived rank position in the population preference distribution. Mathematically, $PN^m(x|\boldsymbol{x})$ is a one-dimensional function that converts the preference ordering acquired from the rank-order decision rule into a real value, given the information that the firm introduced the product line $\boldsymbol{x}$. The magnitude of PN^m is represented by the probability density function of a normal distribution, $\mathcal{N}(\mu, \sigma^2)$. The values of the normal distribution parameters are estimated from $\boldsymbol{s_x^m}$, where μ denotes the mean of $\boldsymbol{s_x^m}$ and σ^2 is an unbiased estimator of the variance. Given the nature of the normal distribution, the bell shape of the distribution implies that middle options are most favourable.

[lxxxviii] Since most of the model settings and assumptions are identical to those of the first model, the repeated part of the model description is just a brief reminder.

3.3.1.1 Explore vs exploit

Although the basic properties of preferences$_U$ and preferences$_N$ are same as before, unlike the first model (which simply sums up the respective influences of these two preferences in the learning stage), the present model separates their influences on choice and assumes that consumers will adopt either an exploration or an exploitation strategy in the choice stage. Denoting the parameter that controls explore-exploit tradeoffs as $\gamma \in [0,1]$, this explore-exploit parameter γ specifies the likelihood of adopting an exploration strategy and is assumed to be constant across menus. In the computer simulation, exploration occurs when $\gamma \geq \varepsilon$, where $\varepsilon \in (0,1]$ is randomly drawn at each round. It is assumed that consumers will always explore in the first round.

3.3.1.2 Explore

In an exploration trial, consumers are in the state of searching (i.e., exploring). The option they decide to try may be influenced their perceived social norms. More specifically, choice is, to a degree Φ, determined in proportion to PN^m of available options[lxxxix] and to degree (1- Φ) suffered determined by random noise[xc]. The parameter $\Phi \in [0,1]$ reflects the extent the which choice is affected by social norms in a marketplace as well as the tendency for consumers to utilise market information when making a

---

[lxxxix] The word "available" here refers to options that have inhibition values less or equal to thresholds.

[xc] In order to achieve choice proportional to PN, the choice sampling is modelled as $\sum\{[\Phi \cdot \frac{\text{cumulative sum of available } PN^m}{\max(\text{cumulative sum of available } PN^m)} + (1-\Phi) \cdot \frac{\text{cumulative sum of available } unif(0,1)}{\max(\text{cumulative sum of available } unif(0,1))}] < \varepsilon\} + 1$, where $\varepsilon \in [0,1]$ is an arbitrary random number, the term "available PN^m" refers to the value of the preference$_N$ on an option that is not inhibited at the time of calculation, and "available *unif* (0,1)" similarly denotes an "available" option's assigned value, which is randomly drawn from a uniform distribution. The whole term $\sum(\cdot) + 1$ gives an index of an "available" option. Note that for available *unif* (0,1), the computer assigns random values to all options first and then uses the ones assigned for "available" options. Due to difference in the scales between PN^m and *unif* (0,1), the summation between these two occurs after rescaling.

decision. The latter also relates to how much consumers trust the preference-relevant information conveyed by market context.

3.3.1.3 Exploit

In contrast, in exploitation trials consumers make a decision by exploiting their past choices. As in the first model, consumers in this such a trial make rank-based inferences about their preferences by computing how frequently an option has been chosen relative to other options in their choice history. Consumer i's context-sensitive inferences about market offerings is referred to as their inferred preference, PI_i. The formula for computing this is::

$$PI_i = \{PI_i(x|\boldsymbol{x}): \mathbb{Z}^3_{\geq} \to \mathbb{R}_+ \left| PI_i(x|\boldsymbol{x}) = \frac{\text{nlower}_x + 0.5 * \text{nequal}_x}{\text{nlower}_x + \text{nhigher}_x + \text{nequal}_x}, \forall x \in \boldsymbol{x} \subseteq X \right\},$$

where $nlower_x$indicates the number of alternatives chosen less frequently than the option *x*, $nhigher_x$refers to the number of alternatives chosen more frequently than *x*, and $nequal_x$ denotes the number of alternatives chosen equally frequently to *x* (*x* itself is not included). Recall that PI_i is defined not in terms of absolute real-world quantities, but in relative rank coordinates. Therefore, consumers can only know that they more frequently chose an option at a given percentile of the product line. They do not know their preferences in terms of the absolute quantities of an attribute *m*, $\boldsymbol{s^m}$.

In addition, in this new model, there is no pre-choice, separate learning stage – consumers are assumed to learn and adjust their inferred preferences every time they make a decision. Finally, the inferred preferences deterministically inform a purchasing decision in exploitation trials, assuming that the income constraint is not binding[xci]. In practice, the choice in such trials is given by $\sum(\frac{\text{cumulative sum of available } PI_i}{\text{max (cumulative sum of available } PI_i)} < \varepsilon) + 1$, where ε is a random number that lies between zero and one and the word "available"

[xci] This presumption rules out any price effects, and thereby allows the analysis to illuminate contextual influences on purely preference-guided consumer behaviour.

in the formula means consumers only take into account options whose inhibition of return is less than their corresponding threshold values .

3.3.1.4 Post-choice stage

In the post-choice stage, regardless of which strategy is employed, for every choice made the chosen option x_j will gain one more point in its choice history, denoted as h_j. Thus, the choice history vector $\boldsymbol{h}_i = [h_{1,i}\ h_{2,i} \ldots\ h_{j,i} \ldots\ h_{n,i}]$ records the total numbers of times that an individual option x_j has been chosen by the consumer i under a certain context. Importantly, due to the fact that memory fades as time passes, at the beginning of each purchasing round, $\boldsymbol{h}_i$ decays[xcii] with a rate of τ. Since PI_i is computed based on this imperfect choice history in memory, the parameter $\tau \in [0,1]$ indirectly controls the relative impact of early and recent past choices on the current purchasing decision. The smaller τ is, the stronger the recency effect is.

3.3.1.5 Rationale for using inhibition of return

Due to the complexity of the method of computation of inhibition of return, it is difficult to introduce the function of inhibition of return and how it operates to prevent an undesirable choice. The description of inhibition of return, thus, begins with rationale of building it. After each choice, consumers may encounter more or less negative feelings towards the chosen product. This may happen because the choice itself is unsatisfactory or simply because consumers are weary of that option. To capture this idea, the model assumes that consumers will inhibit the chosen option and that market options whose levels of inhibition are strictly greater than some threshold values (see the next subsection) are excluded from the consideration set. Conversely, if an option's level of inhibition is less than its threshold value, it will be included. Note that since the idea of inhibition of return is used to capture the fact that consumers often do not consider products known to be unfavourable based on past consumption experiences, it is reasonable to assume that a chosen product will be assigned a high inhibition value if its associated preference$_U$ is low.

[xcii] Choice history at the beginning of each round is choice history in the previous round multiplied by τ.

Mathematically, inhibition of return is computed as an inverse function of chosen products' associated preference$_U$ value. Combined with thresholds, the consequence is that the higher the inhibition is, the more likely it is that the product will not be considered in a given choice round, holding the threshold value the same.

Another important feature of inhibition is that the amount of inhibition is immune from social influence. Social norms only play a role in determining threshold values. This may sound counter-intuitive, but the purpose of modelling in this fashion is to make inhibition more straightforwardly indicative of consumers' feelings towards their choice. That is, inhibition is expected to reflect people's immediate, direct emotional responses to their choice after consumption. It thus should be separated from rational market inferences. By contrast, thresholds often intervene when consumers are faced with a choice task and are deliberately used as a criterion to evaluate an option's eligibility for consideration.

Moreover, according to research studies in psychology, inhibition of return usually spreads to adjacent areas (Birmingham, Visser, Snyder, & Kingstone, 2007; Christie, Hilchey, & Klein, 2013; Taylor, Chan, Bennett, & Pratt, 2015). Thus, it is also assumed that products that are conjectured to be worse than the chosen one will be inhibited following a choice. In other words, if consumers are able to identify whether the chosen product have too much or too little attribute *m* relative to the desired level determined by preferences$_U$, they will assign inhibition to products that are logically considered worse than the chosen one.

The more detailed and formal description of inhibition of return will be presented below. To make it more understandable, it will be helpful to introduce thresholds first because calculation of inhibition involves a use of thresholds.

3.3.1.6 Thresholds

Thresholds are calculated as follows. Recall that unlike in the first model, thresholds in this model work in a reverse way – an option whose inhibition is less than the relevant threshold value will be included in the consideration set. Therefore, counter-intuitively, passing a threshold means that a level of inhibition is less than a threshold value. In addition, the threshold value of each market option in this model will be influenced by social norms if consumers believe that menu contains population's preference-relevant information, i.e., $\Phi > 0$. The threshold equation represents the idea that consumers (with $\Phi > 0$) are more willing to give a future chance of being chosen to options that are perceived to be desirable based on market inferences. In other words, an option is more likely to be in future consideration sets when it is located at middle versus extremes and is in a high social influence setting versus a low one. Mathematically, the threshold value of any market alternative $x \in \boldsymbol{x}$ is expressed as $T(x|\boldsymbol{x}) = t \cdot (1 + \Phi \cdot e^{PN^m(x|\boldsymbol{x})})$. The parameter $t \in \mathbb{R}_+$ is an arbitrary constant that determines the threshold value in the absence of social influence. Note that thresholds depend only on the social norm (and on Φ which is fixed for an individual and context).

3.3.1.7 Inhibition

This subsection will explain how inhibition is calculated and spread to necessarily worse products in the market. The inhibition function is designed to meet the following intuitively plausible requirements:

a) The amount of inhibition calculated using the following formula is an increment to past (accumulated) inhibition after a choice is made.

b) More inhibition should be added to a chosen item if its preference$_U$ is low.

c) Inhibition must always be positive.

d) The formulation should specify the number of time cycles necessary for a product with zero preference$_U$ value to be reconsidered in the choice stage in the no social influence case. Importantly, the calculation does not take inhibition inherited from other chosen options into account. Therefore, it is possible for a product never to be considered again due to continually receiving inhibition from others.

e) The formula of inhibition should specify a level of preference$_U$ such that a chosen product with preference$_U$ below this level should not be considered in next round, no matter what.

f) Non-chosen items should also be inhibited if it is logically obvious that they should be.

g) Inhibition decays over time. Specifically, it decays at the beginning of each round.

Based on the requirements above, the formula of inhibition is developed and presented as follow. First, it is worth emphasising again that the value of inhibition computed using the following formula is an amount added to past inhibition after a choice is made. Second, the functional form of calculating inhibition of return follows a Gaussian function, so as to guarantee only positive inhibition and to distinguish between unfavourable and favourable options. Individual i's inhibition for a chosen $x \in \boldsymbol{x}$, is $Inhibition_i(x) = a \cdot e^{(-b\cdot(PU_i^m(x))^2)}$, where $a = \frac{t}{\tau^z}$ and $b = \frac{-\ln(\tau^{z-1})}{(\max(PU_i^m)\cdot g)^2}$. Consistent with intuition, this has the effect of adding more inhibition to a chosen item if its associated preference$_U$ is small. Note that the input *x* here indeed refers to the choice *x*'s actual attribute value, s_x^m. The whole function $Inhibition_i(x)$ operates on the normalised real attribute space, which is independent of the attribute values offered in the market, of the relative ranks of options, and of choice history. Importantly, the incremental value of inhibition can be computed for any attribute value. The resulting quantity specifies the amount of inhibition that an option with a certain attribute value will accrue if it is chosen. Therefore, the incremental value of inhibition is independent of market offerings or choice history. Moreover, the presence of a preference$_U$ in the formula for inhibition implies that whether consumers are going to make a future purchasing choice and which option should be considered largely depend on the preferences$_U$ of market alternatives. Social norms affect an option's eligibility for consideration only through their influence on threshold values[xciii].

[xciii] This is because inhibition in this model is used to reflect consumers' immediate feelings towards chosen items. The specific reason why inhibition is independent of market influences is addressed in the previous subsection, named "Rationale for using inhibition of return".

The parameters τ, t, z, and g can be viewed as fixed dispositional traits. As noted before, τ refers to a decay rate of memory and t denotes the threshold value in the absence of social influence. The parameter $z \in \mathbb{N}_+$ in the equation indicates the number of rounds required for an option with zero preference$_U$ value to be reconsidered in the choice stage in the no-social-influence case[xciv]. Importantly, since z does not change with Φ, the maximum number of rounds that must pass before zero-preference$_U$ options can be reconsidered in the case where $\Phi > 0$ is expected to be smaller than z itself, thanks to higher threshold values.

A further intuition concerns the need to ensure that options with sufficiently low preferences$_U$ (defined as < b) must not be in the consideration set for the next round[xcv]. Therefore, the denominator of the variable b, (max $(PU_i^m) \cdot g)^2$, gives the smallest value of the preference$_U$ that satisfies $\tau \cdot Inhibition_i \leq t$. For convenience, this lower-bound value is set to be proportional to the maximum preference$_U$ over the technologically feasible set, and the value of the proportion is controlled by the parameter $g \in (0,1] \subseteq \mathbb{R}_+$.

3.3.1.9 Indifference area and spreading of inhibition

It is often observed in reality that consumers can sometimes report their relative preferences between products, but sometimes cannot. This is possibly because if products

---

[xciv] The process of deriving the expression of the variable a in terms of parameters τ, t, and z is shown as follows. First, substitute $PU_i^m(x) = 0$ into the equation of $Inhibition_i(x)$ to get $Inhibition_i(x) = a \cdot e^{(-b \cdot (0)^2)} = a$. Then, by the definition of z, $a \cdot \tau^z \leq t$. Rearranging it, $a \leq \frac{t}{\tau^z}$ is obtained.

[xcv] The following steps were used to derive the variable b with respect to parameters τ, z, and g. Given the definition of g, the inequality $a \cdot e^{(-b \cdot (\max (PU_i^m) \cdot g)^2)} \cdot \tau \leq t$ is obtained. Rearrange the inequality, $e^{(-b \cdot (\max (PU_i^m) \cdot g)^2)} \leq \frac{t}{a \cdot \tau}$. Then, take the ln of both sides and rearrange the inequality, $-b \cdot (\max(PU_i^m) \cdot g)^2 \leq \ln\left(\frac{t}{a \cdot \tau}\right) \Rightarrow b \geq \frac{-\ln\left(\frac{t}{a \cdot \tau}\right)}{\left(\max(PU_i^m) \cdot g\right)^2}$. Last, substitute the equation $a = \frac{t}{\tau^z}$ into the inequality, $b \geq \frac{-\ln\left(\tau^{z-1}\right)}{\left(\max(PU_i^m) \cdot g\right)^2}$.

are satisfactory, people may feel they are indifferent. To incorporate this observation, it is assumed that options with preferences$_U$ greater than the lower bound value (i.e., $\max(PU_i^m) \cdot g$) are all believed to be "good enough". Among these options, the consumer i is assumed not to know which one is preferred to others[xcvi], although their associated amounts of inhibition of return (if chosen) are different[xcvii]. Therefore, a higher value of g reflects a small "indifference area."

Put the other way around, market options lying outside the indifference area are assumed to be consciously distinguishable as their associated preferences$_U$ values are too low to satisfy consumers. The model assumes that consumers are capable of identifying whether the attribute values of these options are too high or too low, relative to the optimum. That is, consumers can clearly recognise that options that are further away are even worse than the chosen one. According to this point of view, it is reasonable to assume that options that are thought to be worse than the choice will inherit the same amount of inhibition of return as the chosen option[xcviii]. Formally, the conditions for any market alternative $x \in \boldsymbol{x}$ to inherit inhibition from the choice $x_j \in \boldsymbol{x}$ are (1) $PU_i^m(x_j) < \max(PU_i^m) \cdot g$, and (2) either $s^{m*} > s_j^m \geq s_x^m$ or $s_x^m \geq s_j^m > s^{m*}$, where s^{m*} refers to

---

[xcvi] This idea is consistent with the assumption of the first model, which suggests that consumers cannot make a pairwise choice when the difference between the two options' preferences$_U$ are less than an arbitrary threshold.

[xcvii] As noted, the magnitude of inhibition of return can be interpreted as the strength of a negative feeling on an option. It is akin to an emotional response based on one's intrinsic preferences, rather than a conscious assessment of an option's badness. Hence, like preferences$_U$, the value of inhibition of return is consciously inaccessible.

[xcviii] The reason why consumers are assumed to assign the same level of inhibition, instead of more, to obviously worse products is because consumers may not know how much worse other products are than the chosen one. They will likely be unaware of actual attribute levels of market alternatives and hence of their real feelings of products that are thought to be worse. Even if consumers have tasted those worse products before, they are likely to forget the exact feelings they had. Therefore, the only sure thing is that the worse products should be get inhibition at least as much as the chosen product.

the optimal level of the attribute *m*. Note that the first condition stresses that the spill-over effect only happens on alternatives outside the indifference area. For those "good enough" options, consumers cannot identify the preference direction, and thus are unable to distribute inhibition to other options.

As with the choice history vector, the size of inhibition, including inhibition generated by the option itself when it is chosen and inhibition inherited from other alternatives in the rounds that others are chosen, will accumulate across iterations but will decay with a constant rate of τ at the beginning of each round. This reflects the fact that bad memories or negative impressions about options will also fade over time.

Finally, the model assumes that if the monopolist changes the product line, the consumer *i* will re-construct the all of the variables, including choice history, inhibition of return, and preferences$_\mathrm{I}$, through the same process under the new context. Past contexts will not influence consumers' evaluations on market offerings in the new context at all.

3.3.2 The firm's profit-maximisation problem and the market equilibrium

The analysis of the monopolist's best response to market demand is identical to the one explored in Chapter 2. This subsection is a summary of its counterpart in that chapter. The aim of the section is to briefly remind readers of how the profit-maximising firm will react to market demand under two cases: (1) products' marginal costs and prices are constant, and (2) products' marginal costs and prices linearly increase with their attribute values.

Let p_j and c_j be exogenously determined prices and marginal costs of the product x_j, respectively. Moreover, a fixed cost *FC* is assumed to exist and to remain unchanged for a certain range of output levels. Then, denote $\pi(x_j; \boldsymbol{x})$ as the profits obtained from selling x_j under the context $\boldsymbol{x}$. The total profit earned by offering $\boldsymbol{x}$ is $\Pi(\boldsymbol{x}) \equiv \sum_{j=1}^{n} \pi(x_j; \boldsymbol{x}) = \sum_{j=1}^{n} [p_j \cdot q_j^D(\boldsymbol{x}) - c_j \cdot q_j^s(\boldsymbol{x})] - FC$, where $q_j^D(\boldsymbol{x})$ and $q_j^s(\boldsymbol{x})$ respectively refer to quantities sold and quantities supplied of the option x_j under the context $\boldsymbol{x}$. Suppose that the monopolist has reliable information about demand for x_j given product line $\boldsymbol{x}$, i.e., $D(x_j; \boldsymbol{x})$ is known. To solve the profit-maximisation problem, the firm will first identify its profit-maximising supply function, $S\big(x_j;\boldsymbol{x}\big|D\big(x_j;\boldsymbol{x}\big)\big)$[xcix], which gives the profit-maximising output level of x_j under each feasible context, and then choose a product line $\boldsymbol{x}$ that will lead to maximum profits. That is, the firm will use information about $D\big(x_j;\boldsymbol{x}\big)$ to determine optimal output (supply) levels for each context, and then choose the context that maximises total profits, given consumer demand and the optimal output levels.

Selection of the profit-maximising context involves determining the number of options, *n*, as well as choosing the attribute values of all products $\boldsymbol{s_x^m}$. To simplify the analysis, it is assumed that the magnitudes of the attribute *m* of each market offering form an arithmetic progression, whereby products' attribute levels differ from their nearby neighbours' by a common difference, $d \in \mathbb{R}_{++}$. Mathematically, $s_{j+1}^m - s_j^m = d \; \forall \; j = 1, \ldots . n-1$.

3.3.2.1 Case 1: Constant price and marginal cost

This subsection analyses the firm's profit-maximisation problem in the case where products' marginal costs and prices are independent of their attribute levels. That is, the price p_j, the marginal cost c_j, and their differences $(p_j - c_j)$ are fixed for all *j*. Taking consumer demand $D\big(x_j;\boldsymbol{x}\big)$ as given, the firm obtains $S\big(x_j;\boldsymbol{x}\big|D\big(x_j;\boldsymbol{x}\big)\big)$ by maximising profits with respect to q_j^S for each possible $\boldsymbol{x}$: $\max_{q_j^S \in \mathbb{R}_+} \Pi \equiv p \cdot \sum_{j=1}^{n} q_j^D(\boldsymbol{x}) - c \cdot \sum_{j=1}^{n} q_j^S(\boldsymbol{x}) - FC$, subject to $q_j^S(\boldsymbol{x}) \geq q_j^D(\boldsymbol{x})$. The constraint exploits the fact that for any market context, quantity supplied of good x_j should be at least as much as its quantity sold – consumers are unable to buy products that have not been produced. To obtain a

---

[xcix] In other words, $S\big(x_j;\boldsymbol{x}\big|D\big(x_j;\boldsymbol{x}\big)\big)$ shows the profit-maximising quantity supplied of x_j in each of all feasible contexts, given the market demand $D\big(x_j;\boldsymbol{x}\big)$. By its nature, it portrays a relationship between the profit-maximising output level of x_j and context $\boldsymbol{x}$.

solution, the first order necessary condition[c] implies that the optimal quantity supplied under a context $\boldsymbol{x}$, namely $q_j^{s*}(\boldsymbol{x})$, should equal the quantity sold, $q_j^D(\boldsymbol{x})$, for any option x_j. This suggests that when the market clears, there is no excess supply for any market offering in the partial equilibrium, i.e., $S(x_j;\boldsymbol{x}|D(x_j;\boldsymbol{x})) = D(x_j;\boldsymbol{x})$.

The second step in solving the problem is to choose the profit-maximising context, $\boldsymbol{x}^*$, such that $\boldsymbol{x}^* \in \underset{\boldsymbol{x}\subseteq X}{\operatorname{argmax}}\,(p-c)\cdot\sum_{j=1}^{n} D(x_j;\boldsymbol{x}) - FC$[ci]. The first order and second order conditions imply respectively that $(p-c)\cdot\frac{\partial\sum_{j=1}^{n} D(x_j;\boldsymbol{x}^*)}{\partial\boldsymbol{x}^*} = 0$ and $(p-c)\cdot\frac{\partial^2\sum_{j=1}^{n} D(x_j;\boldsymbol{x}^*)}{\partial(\boldsymbol{x}^*)^2} < 0$. Specifically, the first order necessary condition suggests that for any $(p-c) \neq 0$, the profit is optimised at a context that either maximises or minimises total demand. The second order condition further indicates that the sign of s$(p-c)$ and $\frac{\partial^2\sum_{j=1}^{n} D(x_j;\boldsymbol{x}^*)}{\partial(\boldsymbol{x}^*)^2}$ are different and that neither of them is zero. This means that if $(p-c) > 0$, the optimal context is the one that maximises market demand. Conversely, if $(p-c) < 0$, the optimal context minimises market demand. Accordingly, the solution is to choose the contexts that maximise total market quantity demanded as long as $(p - c) > 0$. Note that the solution implies that $\boldsymbol{x}^*$ is not necessarily chosen to make $\frac{\partial D(x_j;\boldsymbol{x})}{\partial\boldsymbol{x}} = 0$ for all j. That is, the firm only cares whether products are bought; which products are bought is not of concern. Thus in this case the aim of the firm is not to promote a particular product, but to minimise the probability of choosing nothing. Therefore, the firm tends to choose the context(s) that best satisfies consumers' expressed preferences, instead of striving to manipulate context to a certain direction based on the degree of consumer bias.

---

[c] The Lagrangian of this problem is $\mathcal{L}(q_j^s,\lambda) = p\cdot\sum_{j=1}^{n} q_j^D(\boldsymbol{x}) - c\cdot\sum_{j=1}^{n} q_j^s(\boldsymbol{x}) - FC + \lambda(q_j^s - q_j^D)$. This gives the following Karush-Kuhn-Tucker conditions: (1) $\frac{\partial\mathcal{L}}{\partial q_j^s} = -c + \lambda = 0$, (2) $q_j^s - q_j^D \geq 0$, (3) $\lambda \geq 0$, and (4) $\lambda(q_j^s - q_j^D) = 0$. The solution is $(q_j^{s*},\lambda^*) = (q_j^D, c)$.

[ci] Note that $D(x_j;\boldsymbol{x})$ represents sequences of several discrete data points. To treat it as a differentiable function, it is assumed that $D(x_j;\boldsymbol{x})$ here has been transformed into a polynomial, continuous function by using Newton's divided differences approach.

3.3.2.2 Case 2: Attribute-value dependent price and marginal cost

This subsection explores the case where products' marginal costs and prices linearly and positively depend on their attribute values. Consider an alternative state where selling products with higher values of the attribute *m* is more profitable. In this case, p_j and c_j are no longer constant over all *j*. In effect, their values depend on the attribute level of x_j on the attribute *m*. To avoid trade-offs in the price-attribute space, assume *p* and *c* are a linear function of values of the attribute *m* such that $\frac{dp(s^m)}{ds^m} > 0$ and $\frac{dc(s^m)}{ds^m} > 0$ for all feasible s^m. For simplicity, let $p_j = \delta_p \cdot s_j^m$ and $c_j = \delta_c \cdot s_j^m$, where δ_p and δ_c are strictly positive constants.

The first result of the profit-maximising problem in this case is that the firm's optimal quantity supplied under a context $\boldsymbol{x}$ should equal to the quantity sold for any option x_j, consistent with the previous case. To illustrate this, rewrite the total profit function as $\Pi(\boldsymbol{x}) \equiv \sum_{j=1}^{n} \pi(x_j; \boldsymbol{x}) = \sum_{j=1}^{n}\left[\left(\delta_p \cdot s_j^m\right) \cdot q_j^D(\boldsymbol{x}) - \left(\delta_c \cdot s_j^m\right) \cdot q_j^s(\boldsymbol{x})\right] - FC$. Mathematically equivalently, $\Pi(\boldsymbol{x}) = s_1^m \cdot \left[\delta_p \cdot q_1^D(\boldsymbol{x}) - \delta_c \cdot q_1^s(\boldsymbol{x})\right] + (d + s_1^m) \cdot \left[\delta_p \cdot q_2^D(\boldsymbol{x}) - \delta_c \cdot q_2^s(\boldsymbol{x})\right] + \cdots + [(n-1)d + s_1^m] \cdot \left[\delta_p \cdot q_n^D(\boldsymbol{x}) - \delta_c \cdot q_n^s(\boldsymbol{x})\right] - FC$. Given knowledge of consumer demand $D(x_j; \boldsymbol{x})$, the monopolist first derives $S(x_j; \boldsymbol{x} | D(x_j; \boldsymbol{x}))$ by maximising profits with respect to q_j^s for each possible $\boldsymbol{x}$, subject to $q_j^s(\boldsymbol{x}) \geq q_j^D(\boldsymbol{x})$. The solution suggests that the optimal output level, $q_j^{s^*}(\boldsymbol{x})$, that constitutes $S(x_j; \boldsymbol{x} | D(x_j; \boldsymbol{x}))$ is same as previously: $q_j^{s^*}(\boldsymbol{x}) = q_j^D(\boldsymbol{x})$. This implies that no excess supply exists in market equilibrium.

Secondly, it is found that profit-maximising firms will under plausible assumptions choose the context that maximises the "sum of attribute value" sold. With complete information on $q_j^{s^*}(\boldsymbol{x})$ and $q_j^{D^*}(\boldsymbol{x})$ for all possible $\boldsymbol{x}$, the firm will then choose an optimal context $\boldsymbol{x}^*$ to maximise its profits. Mathematically, it will choose $\boldsymbol{x}^* \in \underset{x \subseteq X}{\operatorname{argmax}} \left(\delta_p - \delta_c\right) \cdot \sum_{j=1}^{n}\left[s_j^m(\boldsymbol{x}) \cdot D(x_j; \boldsymbol{x})\right] - FC$. By the same token, the first and second order conditions imply that $\left(\delta_p - \delta_c\right) \cdot \frac{\partial \sum_{j=1}^{n}\left[s_j^m(\boldsymbol{x}^*) \cdot D(x_j; \boldsymbol{x}^*)\right]}{\partial \boldsymbol{x}^*} = 0$ and $\left(\delta_p - \delta_c\right) \cdot \frac{\partial^2 \sum_{j=1}^{n}\left[s_j^m(\boldsymbol{x}^*) \cdot D(x_j; \boldsymbol{x}^*)\right]}{\partial (\boldsymbol{x}^*)^2} < 0$, respectively. This clearly indicates that the attribute values of products in this case have a direct effect on the choice of context. The solution is to choose the contexts that maximise total attribute values sold as long as $(\delta_p - \delta_c) > 0$.

The solution can be illustrated by a simple example. Suppose that for any two possible contexts $\boldsymbol{x_{old}}$ and $\boldsymbol{x_{new}}$, where $\boldsymbol{s^m_{x_{new}}} = \boldsymbol{s^m_{x_{old}}} + \varepsilon$, $\varepsilon > 0$, the quantity demanded of options with the same relative ranks under $\boldsymbol{x_{old}}$ and $\boldsymbol{x_{new}}$ is identical at market equilibrium, i.e., $q_j^{D^*}(\boldsymbol{x_{old}}) = q_j^{D^*}(\boldsymbol{x_{new}}) \forall j \in \{1, \dots, n\}$. The firm's best strategy will be choosing the context with higher attribute values, namely $\boldsymbol{x_{new}}$, over $\boldsymbol{x_{old}}$. This further implies that if by any chance, consumer demand on each market alternative x_j is identical across contexts and at least one $q_j^{D^*} > 0$, i.e., the curve of $D(x_j; \boldsymbol{x})$ is a horizontal line above zero, the firm will be better off by setting $\boldsymbol{s^m_x}$ as high as possible as long as $(\delta_p - \delta_c)$ and there will be no limit to this process. In other words, the equilibrium context will be as close as possible to the upper limits of the technologically feasible set of an attribute *m*.

However, consumer behaviour in the proposed model is of course more complex than in the above example, and consequently $q_j^{D^*}(\boldsymbol{x})$ is not always the same for all *j* across contexts. In effect, once s_j^m is beyond the level that generates the maximum preference$_U$ value, any increment in s_j^m is likely to be accompanied by a loss in q_j^D, which yields a trade-off in determining $\boldsymbol{s^m_x}$. Yet, for consumers who are strongly biased towards choosing middle option(s) owing to the false consensus effect, their quantity demanded for middle options will be hard to change, unless the increment in s_j^m is too large. This gives the firm an incentive to distort context. Therefore, it is expected that $\boldsymbol{x}^*$ here may be closer to the upper limit of the attribute space than in the first class of the profit-maximisation problem (i.e., when profit is independent of attribute value), since the presence of s_j^m in the first and second order conditions allows profit-maximising context to (slightly) move away from the preference$_U$-maximising point. In this case, consumers who suffer a stronger degree of social influence may be exploited more than those who are relatively less biased, since the false consensus effect will predispose biased consumers to choose the middle option(s) when options with smaller values of attribute *m* are indeed preferred in terms of preferences$_U$.

3.3.2.3 Summary

This formulation of the model suggests that when marginal costs and prices are constant, the firm does best if it offers a context that maximises total market quantity demanded, as long as prices exceed marginal costs. That is, the distribution of market demand over market alternatives does not affect profit. Therefore, there is no incentive for the monopolist to exploit biased consumers by distorting context. In contrast, results from the second case, where products' marginal costs and prices are linearly dependent on attribute values, suggest that the attribute values of products have a direct effect on the choice of context. The firm here is best off choosing contexts (i.e., sets of market options) that maximise total attribute values sold. The firm hence has a monetary incentive to manipulate the product line to increase the quantity demanded for more profitable products.

3.4 Computer simulation of the compromise effect

3.4.1 Preliminaries

This section uses computer simulations to illustrate the model's predictions of human choice behaviour under various contexts. In particular, the aim is to show how biased market inferences contribute to the existence of compromise behaviour in equilibrium. The firm's responses in changing contexts and consumer welfare in market equilibrium will also be explored. As in Chapter 2, the simulation will begin with a simple single agent case and then move to the multi-agent case to simulate a real market scenario. Also, for simplicity and transparency, throughout all simulations, the number of market offerings is fixed at five and all offerings are ordered along one dimension. Importantly, it is assumed that consumers who make no choice in the market will choose an outside option x_0 instead. The simulation is programmed with Matlab 2020b.

3.4.1.1 A rational choice benchmark model

To examine the extent of deviation from rational choice and the degree of the compromise effect, a benchmark model that operates within a framework of logit choice rules (e.g., Fudenberg & Kreps, 1993; Luce, 1959) will be used in the simulation. That is, a softmax function will be employed as a benchmark to approximate rational/optimal behaviour with random noise. Mathematically, for any option $x \in \boldsymbol{x} = \{x_1, \dots, x_n\}$, its market share (*MS*) obtained from a consumer *i* in the benchmark model is: $MS_i^R(x|\boldsymbol{x}) = \frac{e^{PU_i^m(x)}}{\sum_{j=1}^{n} e^{PU_i^m(x_j)}}$. The use of a stochastic choice rule allows the benchmark results to reflect preference orderings of all options, which the conventional utility maximisation theory, characterised by the "winner takes all" rule, cannot have. Obtaining information about relative preferences of options is of crucial importance in computing the degree of the compromise effect since it often requires market shares of non-dominant options as inputs. That being said, preference$_U$-maximising choice will still be computed and presented as an additional rational choice benchmark. Finally, since there seems no sensible way to assume a certain choice threshold level for rational/ benchmark agents, profit-maximising contexts as well as equilibrium welfare predicted by the proposed models will not be compared to the benchmark.

3.4.1.2 Measurement approaches of the compromise effect

According to Neumann, Böckenholt and Sinha (2016), there are three commonly-adopted methods of quantifying context effects: Middle proportions, absolute-share changes, and relative-share changes. Yet, none of these specifies the strength of the effect[cii]. This motivates the present paper to develop three different formulae to compute the degree of the compromise effect based upon each of these three measurement approaches. All these three formulae use market share rather than choice probability, because benchmark choice probability is unknown. Note that given that the number of market alternatives is fixed at five for all possible contexts in the computer simulations presented in this section, the analysis of contextual influence is based on a quinary-quinary set comparison. Consistently, for purpose of clarity, the following illustration of

[cii] The traditional methods only show how much (absolute/ relative) choice probability changes due to contextual change, without specifying how strong the change is on a normalised 0-to-1 scale. This makes it difficult to compare the degree of the compromise effect across contexts and models.

each approach is based on a situation where the firm offers five products in the market. Therefore, x_3 refers to the middle option, whereas x_1 and x_5 are the most extreme options in any menu.

3.4.1.3 Middle Proportions

The first assessment approach uses the classical method called "middle proportions", which treats the share of choices taken by the middle option as an indicator of the propensity to compromise (Simonson & Nowlis, 2000). The advantage of middle proportions is that it is straightforward and results calculated using it are easy to interpret and explain. This method, however, does not capture the possibility that a high choice share of a middle option is a result of its high associated preferences$_U$. In other words, the option may be attractive in terms of preferences$_U$, independently of its position in a choice set. To precisely compute the degree of the compromise effect and meanwhile retain benefits of using middle proportions as a measure, a formula was developed based on the original middle proportions approach.

The new formula was created based on the idea that by comparing the proposed model's estimate of market share of the middle option (x_3), noted as $MS_i^P(x_3|\boldsymbol{x})$, with the benchmark result, namely $MS_i^R(x_3|\boldsymbol{x})$, it is possible to cancel out the share gained from the attractiveness of the middle option itself. Moreover, to derive the strength of contextual influences for each possible market context, the (log) difference in estimated market share of a middle option between the proposed model and the benchmark is compared with the (log) difference between benchmark market share of the middle option and 1. The number, 1, here represents the maximum possible market share the middle option can gain. The (log) difference between benchmark results and 1 thus represents how much *additional* market can accrue to the middle option due to the compromise effect when the compromise effect is at its maximum possible. Accordingly, by using the new formula of middle proportions, the maximum degree of the compromise effect is observed when market share of the middle option estimated by the proposed model is 1 and is larger than the benchmark.

The degree of the compromise effect is computed as $\frac{\ln\left(MS_i^R(x_3|\boldsymbol{x})\right)-\ln\left(MS_i^P(x_3|\boldsymbol{x})\right)}{\ln\left(MS_i^R(x_3|\boldsymbol{x})\right)}$, which is a simplified version of a more intuitively explicable formula [ciii], $\frac{\ln\left(MS_i^P(x_3|\boldsymbol{x})/MS_i^R(x_3|\boldsymbol{x})\right)}{\ln\left(1/MS_i^R(x_3|\boldsymbol{x})\right)}$. The natural logarithm is used to indicate the "growth rate", which gives the percentage change in the share of the middle option treating the benchmark result as the baseline. Therefore, the numerator specifies the percentage by which the share of the middle option estimated by the proposed model is different from the benchmark. In other words, this implies the likelihood that consumers in the proposed model chooses a middle option from the choice set more or less frequently than the benchmark consumers do. The denominator, in similar vein, represents the situation where the compromise effect is strongest (in which the share is raised from its baseline value to 1). Thus with this formula, the maximum predicted degree of contextual influences is normalised to 1.

However, this approach is incapable of estimating the degree of any effect in the reverse direction. If the estimated share of the middle option is smaller in the proposed model, i.e., $MS_i^P(x_3|\boldsymbol{x}) < MS_i^R(x_3|\boldsymbol{x})$, the predicted effect size will turn negative and, more problematically, it may easily exceed -1. In effect, by setting the target share to be 1 in the denominator, the formula can necessarily only account for only one direction of contextual effect. Thus, in all of the following simulations, all negative predicted sizes will be set as zero and interpreted as if there is no compromise effect. Furthermore, using

[ciii] The term $\frac{\ln\left(MS_i^P(x_3|\boldsymbol{x})/MS_i^R(x_3|\boldsymbol{x})\right)}{\ln\left(1/MS_i^R(x_3|\boldsymbol{x})\right)}$ is mathematically equivalently to $\frac{\ln\left(MS_i^P(x_3|\boldsymbol{x})\right)-\ln\left(MS_i^R(x_3|\boldsymbol{x})\right)}{\ln(1)-\ln\left(MS_i^R(x_3|\boldsymbol{x})\right)}$, according to the logarithm quotient rule. Given the fact that $\ln(1) = 0$, the terms can be further simplified to be $\frac{\ln\left(MS_i^P(x_3|\boldsymbol{x})\right)-\ln\left(MS_i^R(x_3|\boldsymbol{x})\right)}{-\ln\left(MS_i^R(x_3|\boldsymbol{x})\right)}$, which is equivalent to $\frac{-\left(\ln\left(MS_i^P(x_3|\boldsymbol{x})\right)-\ln\left(MS_i^R(x_3|\boldsymbol{x})\right)\right)}{\ln\left(MS_i^R(x_3|\boldsymbol{x})\right)}$. According to the distributive rule, it is equal to $\frac{\ln\left(MS_i^R(x_3|\boldsymbol{x})\right)-\ln\left(MS_i^P(x_3|\boldsymbol{x})\right)}{\ln\left(MS_i^R(x_3|\boldsymbol{x})\right)}$.

the softmax function as the benchmark model is essential in this case as it guarantees that the denominator will be non-zero.

3.3.1.4 Absolute-share changes

The second measure is based on changes in absolute share, and was developed by Simonson (1989). This paradigm examines the compromise effect by comparing the market share of a certain option before and after it becomes the middle option. It therefore complements the first measure as it directly measures the effect of contextual change on an option and tests the regularity principle. Yet, without contrasting with the benchmark prediction, the result derived from this measure may be misleading. This is because when choice involves a stochastic element, changes in absolute share can be quite natural – the removal of the most attractive option, e.g., the first option, can effectively increase the popularity of all remaining options. This motivates the present paper to develop a new measure to assess the degree of the compromise effect based on changes in absolute share.

Due to the complexity of the new formula, an explanation of the rationale used for its development may be useful. Firstly, the term "absolute-share changes" in the formula is computed as the percentage change in absolute market share of an option before and after contextual change. In practice, the conventional method of computing a growth rate is employed to calculate a percentage change in market share. That is, the difference between the natural logarithms of an option's market share before and after contextual change is used in the formula. Secondly, since the formula was developed to measure the strength of the compromise effect, the option of interest is the one that is not a middle option in the initial context but becomes the middle one after a contextual change. Therefore, the formula focuses on this option's percentage change in market share before and after contextual change. For convenience, all options mentioned in the following refers to this type of option.

As noted in a previous subsection, if the strength of the compromise effect is at its maximum, market share of the option will be raised to 1 after contextual changes, no matter what its market share in the original context is. Therefore, the maximum possible value of the absolute-share change of the option is computed as the difference between the natural logarithms of the option's market share before the contextual change and 1. To obtain the degree of the compromise tendency for consumers in the proposed model, the percentage change in market share of the option before and after a contextual change estimated in the proposed model is divided by the maximum possible value stated above. This measure corresponds to the first term (i.e., the first $[\cdot]$) of the numerator of the formula.

However, this is not the end of the story. As stated, an observed increase in market share of the option after a contextual change can occur simply because the most attractive option is removed from the set. Accordingly, the degree of the "compromise tendency"[civ] for benchmark consumers must also be computed to serve as a reference point. Similarly to what has been done in the proposed model, the benchmark compromise amount is computed by dividing the estimated percentage change in the benchmark market share of the option before and after a contextual change by the maximum possible value. This measure is shown as the second term (i.e., the second $[\cdot]$) of the numerator of the formula. The difference between the degree of the compromise tendency estimated in the proposed model and the benchmark reflects how much consumers in the proposed model are biased towards an item placed at the middle in the set.

Yet, simply taking the absolute difference between the degrees cannot accurately give the strength of the compromise effect. One primary reason is that the absolute difference largely depends on the size of the benchmark result. For example, suppose that the benchmark market share of an option increases from 0.5 to 0.9 after a contextual change. The estimated degree of the compromise tendency in the benchmark model will be $\frac{\ln(0.9)-\ln(0.5)}{\ln(1)-\ln(0.5)} = 0.8481$. Such a high value leaves little room for increase. In other words, even if consumers in the proposed model exhibit the strongest possible compromise effect, the estimated of it derived from taking the absolute difference is just $1 - 0.8481 = 0.1519$.

[civ] Although the term compromise is used, it does not mean benchmark consumers have a preference towards a middle option. As noted before, an item's relative position in the set does not affect benchmark choice. The reason of using the word "compromise" is to emphasise that the number to compute is the benchmark counterpart of the degree of the compromise tendency calculated in the proposed model.

To resolve this problem, the absolute difference between the two compromise effect measures is divided by the absolute difference between the benchmark degree of the compromise tendency and the degree computed in the proposed model, with the assumption that consumers in the proposed model exhibit the strongest compromise effect (i.e., their degree of the compromise tendency is 1). By doing so, the formula can produce a more precise estimate of the actual degree of the compromise effect. Continuing the previous example, the final estimated degree of the compromise effect now is $\frac{1-\frac{\ln(0.9)-\ln(0.5)}{\ln(1)-\ln(0.5)}}{1-\frac{\ln(0.9)-\ln(0.5)}{\ln(1)-\ln(0.5)}} = \frac{1-0.8481}{1-0.8481} = 1.$

In summary, the value computed using this formula gives the strength of the compromise effect in terms of absolute-share changes. It is worth noting that if percentage changes in market share of the option is purely a result of a change in available products' associated preferences$_{U}$ (i.e., if no systematic bias is involved), the value computed with the formula will be zero. Moreover, just as with the middle proportions approach, this formula only works when becoming a middle option can increase an option's market share. If the market share of an option decreases after it becomes the middle option of the choice set, the estimated effect size computed using the formula will be misleading. In addition, if the degree of the compromise tendency estimated by the proposed model is smaller than the benchmark result, the final value produced by the formula is meaningless too. This is because in such a case consumers in the proposed model may less prefer options placed in the middle of the set than other options, but the formula cannot precisely indicate the size of the reversed compromise effect.

The next part of this subsection shows the formula. As before, the illustration of the formula uses x_3 to represent the middle option in the choice set. The notations $\boldsymbol{x_1}$ and $\boldsymbol{x_2}$ are used to denote the contexts before and after change. More specifically, the change in contexts involves dropping an option with the lowest attribute value in the original context, i.e., dropping x_1 in $\boldsymbol{x_1}$, and adding an option with an attribute value greater than any pre-existing alternative, i.e., adding x_5 to $\boldsymbol{x_2}$. As the result, the original 5th option in $\boldsymbol{x_1}$ is now x_4 in $\boldsymbol{x_2}$. Furthermore, since the formula is developed to measure the strength of the compromise effect, the option used in the illustration is x_4 in $\boldsymbol{x_1}$ (and then it becomes x_3 in $\boldsymbol{x_2}$)[cv].

The effect size is computed as follow:

$$\frac{\left[\frac{\ln\left(MS_i^P(x_4|\boldsymbol{x_1})\right)-\ln\left(MS_i^P(x_3|\boldsymbol{x_2})\right)}{\ln\left(MS_i^P(x_4|\boldsymbol{x_1})\right)}\right]-\left[\frac{\ln\left(MS_i^R(x_4|\boldsymbol{x_1})\right)-\ln\left(MS_i^R(x_3|\boldsymbol{x_2})\right)}{\ln\left(MS_i^R(x_4|\boldsymbol{x_1})\right)}\right]}{1-\left[\frac{\ln\left(MS_i^R(x_4|\boldsymbol{x_1})\right)-\ln\left(MS_i^R(x_3|\boldsymbol{x_2})\right)}{\ln\left(MS_i^R(x_4|\boldsymbol{x_1})\right)}\right]}.$$

The rationale behind this formula is consistent with the one using the middle-proportion approach. Each numerator component [·] represents the percentage change in the absolute share of a particular option before and after contextual change. In other words, it specifies how much additional share an option can obtain or lose when it becomes the middle option, x_3. The number 1 in the denominator instantiates the idea that the strongest compromise effect is defined as the one yielding $MS_i^P(x_3|\boldsymbol{x_2}) = 1$, given $MS_i^R(x_3|\boldsymbol{x_2}) < 1$. Therefore, the denominator returns the value that the strongest compromise effect should produce.

Finally, as stated, this measure only captures the compromise effect in one direction. In following simulations, any negative effect size will therefore be turned into zero. Moreover, by the property of logarithms, for the formula to work all inputs, i.e., market shares, are assumed to be non-zero[cvi] and share in the first context, i.e., $MS_i^P(x_4|\boldsymbol{x_1})$ and $MS_i^R(x_4|\boldsymbol{x_1})$, must be strictly below 1. In addition, the condition $\ln\left(MS_i^R(x_4|\boldsymbol{x_1})\right) \neq$

[cv] Although x_4 is used as an example, it is also possible to compute an effect size for other options with the formula. For instance, the estimated degree of the compromise effect for x_5 in $\boldsymbol{x_1}$ can be computed by replacing the notation x_4 to x_5 and $\boldsymbol{x_2}$ to the context where x_5 becomes the middle option.

[cvi] This assumption is likely to be violated in practice. Some options, especially the very unfavourable ones, may generate zero demand and thus zero market share. To avoid an undefined result, when computing via this method the market share of the option of interest will be temporarily set as 0.00001 if it is zero.

$\ln\left(MS_i^R(x_3|\boldsymbol{x_2})\right)$ needs to be met to prevent the denominator from being zero. This implies that the preference$_U$ of the newly introduced option should be different from that of the removed option.

3.3.1.5 Relative-share changes

The third method to investigate the size of the compromise effect incorporates the idea of changes in relative share, as proposed by Simonson and Tversky (1992). This measurement method assesses changes in the relative shares of options in the same choice pair before and after context shifts. This approach is therefore highly relevant to the issue of choice reversal, which occurs when an option is chosen more frequently than another in one context, but less frequently in another. Also, it is useful in detecting violations of the IIA assumption.

To examine the degree of choice reversal and the role of contextual influence, in the following simulations this measurement method will be broken down into two parts, with the first illustrating the probability of choice inconsistency between two contexts and the second computing how much (in %) observed choice inconsistency stems from the compromise effect. Note that the second part is essential because choice inconsistency may just arise from random noise in decision making. For example, for consumers who choose randomly in every context, the probability of choice inconsistency measured between any two context (e.g., the probability of choosing product A over B in one context, but choosing product B over A in another) will be 50%, which is low. Yet, the underlying factor that causes this inconsistency has nothing to do with the compromise effect or any systematic bias. Therefore, the second part is needed to detect the proportion of inconsistency that results from the compromise effect.

For simplicity, the investigation of choice inconsistency focuses on the relative share of x_3 and x_4 in the first context ($\boldsymbol{x_1}$) and x_2 and x_3 in the second context ($\boldsymbol{x_2}$), in which $\{x_3, x_4\}$ and $\{x_2, x_3\}$ are the same choice pair in terms of attribute value. The reason for selecting this pair is that it contains a middle option in each context, allowing the analysis to measure the relative probability of consistently choosing the middle option over its paired option in both contexts.

Choice reversal occurs when consumers choose x_4 over x_3 in $\boldsymbol{x_1}$ while choosing x_2 in $\boldsymbol{x_2}$ or alternatively choose x_3 over x_4 in $\boldsymbol{x_1}$ but choose x_3 over x_2 in $\boldsymbol{x_2}$. Importantly, since choice reversal is defined as inconsistency in choice over the same pair of options before and after context changes, the formula measuring it uses the relative choice frequency of one option over another, rather than the absolute market share of each paired option. Thus, despite the existence of other options in the set (e.g., x_1 in $\boldsymbol{x_1}$, etc.), only the relative market share of the pair $\{x_3, x_4\}$ in $\boldsymbol{x_1}$ and $\{x_2, x_3\}$in $\boldsymbol{x_2}$ is considered. Moreover, although the term "market share" is used, the formula measures individual consumer's degree of choice inconsistency between two contexts, instead of in the market as a whole.

In simulations, the probability of choice reversal for a consumer *i* is computed with the formula:

$$MS_i(x_4|\boldsymbol{x_1};\{x_3,x_4\}) \cdot MS_i(x_2|\boldsymbol{x_2};\{x_2,x_3\}) + MS_i(x_3|\boldsymbol{x_1};\{x_3,x_4\}) \cdot MS_i(x_3|\boldsymbol{x_2};\{x_2,x_3\}).$$

The first multiplication describes the joint probability of choosing x_4 over x_3 in $\boldsymbol{x_1}$ and x_2 over x_3 in $\boldsymbol{x_2}$. Similarly, the second multiplication indicates the joint probability of choosing x_3 over x_4 in $\boldsymbol{x_1}$ and x_3 over x_2 in $\boldsymbol{x_2}$. Both capture the situation of choice inconsistency between two contexts, while only the second multiplication captures a preference towards an option located at the middle of the choice set. Notice that the use of multiplication implies that relative choice frequencies before and after context change are treated as independent events. That is, the relative choice probability in $\boldsymbol{x_1}$ does not affect that in $\boldsymbol{x_2}$.

Nonetheless, as noted, simply assessing probability of choice inconsistency is not sufficient to demonstrate compromise behaviour. Choice reversal may stem from a psychological indifference between options, due to highly similar preference$_U$ values. This indistinguishability is likely to allow an influence of random noise in decision making, which may make the benchmark results more inconsistent than the proposed model's predictions. Therefore, it is necessary to conduct a more nuanced investigation into how much choice reversal is triggered by the propensity to choose the middle options in whatever context. Accordingly, the second formula captures how much the probability of choice reversal computed using the first formula stems from the relative probability of consistently choosing the middle option over its paired option in both contexts, which is exactly the second part of the first context. In other words, the second formula measures the proportion of joint probability of choosing x_3 in both contexts in the probability of choice reversal. In practice, it is computed as follow:

$$\frac{MS_i(x_3|\boldsymbol{x_1};\{x_3,x_4\})\cdot MS_i(x_3|\boldsymbol{x_2};\{x_2,x_3\})}{MS_i(x_4|\boldsymbol{x_1};\{x_3,x_4\})\cdot MS_i(x_2|\boldsymbol{x_2};\{x_2,x_3\})+MS_i(x_3|\boldsymbol{x_1};\{x_3,x_4\})\cdot MS_i(x_3|\boldsymbol{x_2};\{x_2,x_3\})}$$

The formula incorporates the idea that the higher the predicted value is, the more likely it is that the inconsistency has arisen from the compromise effect. Put differently, if this second formula returns a low value, observed choice inconsistency is more likely to reflect random noise, rather than any bias towards the middle option.

3.3.1.6 Consumer welfare

This subsection explains how consumer welfare at market equilibrium is computed. For this welfare analysis, the values of preferences$_U$ will be given a hedonic interpretation. The simulations will first examine average happiness gained from consuming in the market. This is achieved by summing up the product of each option's associated preference$_U$ and its absolute market share in a context. The average welfare gain or loss will be obtained by comparing the proposed model and the benchmark as well as comparing manipulated and non-manipulated contexts. The former reveals the extent to which cognitive constraints and social influences affect welfare gained in the market. The latter addresses the welfare impact of firm's manipulation of context. Furthermore, to capture the welfare effect of the outside option, a condition on utility of an outside option for equalising consumer welfare under manipulated and non-manipulated equilibrium contexts will also be proposed.

3.3.1.7 Notations

Before presenting the simulation results, it is helpful to remind the reader of the definitions of important notation. Table 1 summarises the model's frequently-used notation relevant to all simulations in this section. Note that α_i^m, β_i^m, γ, Φ, t, τ, z, and g are free parameters that all influence consumer i's final choice according to the model.

Table 1. Notation used in the proposed model

Notations	*Definition*
Market	
n	The number of market offerings
x	Any market alternative
x_0	An outside option, i.e., an option that is not offered in the market.
$\boldsymbol{X}$ or $\{x_1, x_2, \ldots, x_n\}$	A set of n market alternatives ordered in terms of magnitudes of a common attribute m
X	A technological feasible set that contains $\boldsymbol{x}$
s_j^m	A normalised attribute value of an option x_j on an attribute m
$\boldsymbol{s_x^m}$	A set of a normalised attribute values of each market alternatives, x, on an attribute m (i.e., a set of s_j^m)
d	A difference in attribute levels between adjacent products
Preferences	
PU_i^m	Consumer i's preferences$_U$ for absolute values of an attribute m
$PN^m(x\|\boldsymbol{x})$	A preference$_N$ of any option x in a context $\boldsymbol{x}$ for an attribute m
$PI_i(x\|\boldsymbol{x})$	The consumer i's preferences$_I$ of any option x in a context $\boldsymbol{x}$
α_i^m, β_i^m	Parameters of probability density function of a beta distribution that represents the consumer i's preferences$_U$ for an attribute m
The proposed model	
γ	The probability of adopting an exploration strategy versus exploitation
Φ	The tendency for consumers to use the rank order decision rule
$h_{j,i}$	The consumer i's choice history over a market alternative, x_j

$\boldsymbol{h}_i$ or $[h_{1,i}\ h_{2,i} \dots h_{j,i} \dots h_{n,i}]$	A set of the consumer i's choice history over each market alternative
$T(x\|\boldsymbol{x})$	A threshold value of any market alternative x in context $\boldsymbol{x}$
t	Constant threshold value in the absence of social influence
τ	Decay rate of memory
z	The number of rounds necessary for an option with zero preference$_U$ value to be reconsidered in the choice stage in the no social influence case
g	A parameter used to specify the lower bound value of an indifference area, i.e., max $(PU_i^m) \cdot g$

3.4.2 Simulation 1: A simple illustration of the compromise effect and equilibrium contexts

The main purpose of this simulation is to demonstrate the presence of the compromise effect and its potential influence on long-run market equilibrium. Consider a situation with a unit-demand agent and two market contexts $\boldsymbol{x_1}$ = {x_1, x_2, x_3, x_4, x_5} and $\boldsymbol{x_2}$ = {x_2, x_3, x_4, x_5, x_6}[cvii], where all market alternatives are arranged in order along their values of the attribute m. Let their corresponding attribute values be $\boldsymbol{s^m_{x_1}}$ = {0.40, 0.50, 0.60, 0.70, 0.80} and $\boldsymbol{s^m_{x_2}}$ = {0.50, 0.60, 0.70, 0.80, 0.90}.

Suppose that the representative consumer i is endowed with a stable preference$_U$, PU_i^m = Beta (5, 15). As revealed by Figure 11 (a), x_1 and x_2 will then be preferred to all other available alternatives under $\boldsymbol{x_1}$ and $\boldsymbol{x_2}$ respectively, since $PU_i^m(s_j^m)$ declines with j for all j = 1, 2, …, 6. As a result, the preference$_U$-maximising choice is {x_1, x_0} in $\boldsymbol{x_1}$ and {x_2, x_0} in $\boldsymbol{x_2}$, depending on the attractiveness of the outside option x_0, since it is assumed that consumers who do not choose in the market will choose x_0. Less strictly, the

---

[cvii] The use of $\boldsymbol{x_2}$ = {x_2, x_3, x_4, x_5, x_6} is to emphasise that the jth option in $\boldsymbol{x_1}$ is the (j-1)th option in $\boldsymbol{x_2}$, in terms of their attribute values. This is to make the explanation of the compromise effect clearer.

benchmark model, computed by the softmax function, suggests the market share of options in $\boldsymbol{x_1}$ will be {0.4292, 0.1669, 0.1364, 0.1338, 0.1337} and {0.2369, 0.1937, 0.1899, 0.1898, 0.1898} in $\boldsymbol{x_2}$. Note that this result does not take x_0 into account, as there is no sensible way to assume a certain choice threshold level for benchmark agents. As stressed previously, the benchmark model cannot predict the probability of not choosing in the market.

However, the absolute magnitude of preferences$_{\mathrm{U}}$ is assumed not to be accessible. This causes choice to be partly shaped by market inferences. At the choice stage, consumers will either explore (i.e., choose based on preferences$_{\mathrm{N}}$) or exploit (i.e., choose based on preferences$_{\mathrm{I}}$). The tendency to explore is controlled by the parameter γ, which was set to be 0.5 in the simulation. Moreover, by fitting a normal distribution to each context $\boldsymbol{x}$, $PN_{\boldsymbol{x_1}}{}^{m} = PN_{\boldsymbol{x_2}}{}^{m} =$ {1.1337, 2.0658, 2.5231, 2.0658, 1.1337} is obtained.

3.4.2.1 Inhibition of return: Pre-iteration predictions

Given parameters τ = 0.9, t = 1, z = 20, and g = 0.7, the expected size of the inhibition of return for every attribute value[cviii] is shown in Figure 11 (b). Notice that the "size" shown in Figure 11 (b) refers to the incremental amount of inhibition that will be added to some products'[cix] total inhibition if a choice is made. In other words, the quantities of inhibition provided in Figure 11 (b) show the additional amount of inhibition that an option with a certain attribute value will receive if it is chosen. This value can therefore be computed for any attribute value, no matter whether these attribute values are offered in the market or chosen by the consumer i.

Moreover, the inhibition curve (Figure 11 (b)) and the threshold curve (Figure 11 (c)) jointly imply that once chosen, the market alternatives, except for x_1 (the first option

---

[cviii] To recap, it is computed using the formula $Inhibition_i(x) = a \cdot e^{(-b \cdot (PU_i^m(x))^2)}$, where $a = \frac{t}{\tau^z}$ and $b = \frac{-\ln(\tau^{z-1})}{(\max(PU_i^m) \cdot g)^2}$. This formula does not require any information about market offerings and choice history.

[cix] The term "some products" refers to a chosen product and possibly products that inherit inhibition from the chosen one.

in $\boldsymbol{x_1}$), will be excluded from the consideration set for nearly 20 purchasing periods when Φ = 0. In other words, it requires nearly twenty rounds for an option's accumulated inhibition to reduce to a value below the constant threshold, *t*. Yet, this value (i.e., 20 purchasing periods) takes neither the spill-over effect nor inherited inhibition values into account. It only uses the expected amount of inhibition immediately generated by a chosen option.

In fact, since contexts $\boldsymbol{x_1}$ and $\boldsymbol{x_2}$ do not contain any "good enough" option, i.e., $\forall j = 1,2,\dots,6, PU_i^m(s_j^m) < 0.7 \cdot \max(PU_i^m) = 2.9424$, existing market alternatives that are, in the attribute space, more extreme than a chosen option will inherit a large amount of inhibition from it. Therefore, it is predicted that products that are not the most preferred ones in the context (according to preferences$_U$), will not be reconsidered for more than 20 purchasing periods when Φ = 0. For the case of Φ > 0, the number of periods that must pass before reconsideration is possible may be fewer than in the case of Φ = 0, because threshold values are increased by preferences$_N$ (see Figure 11 (c)).

3.4.2.2 Probability of not choosing in the market

This subsection investigates the probability of not making any choice from the market for different values of Φ, namely, a tendency to adopt the rank order decision rule, as well as possible factors related to this probability. The selected values of Φ are {0, 0.5, 1}, represent no, medium and high tendency to use market inferences during decision making. The simulation results indicate that over 10000 purchasing periods[cx], the average probability of choosing nothing in the market when Φ = {0, 0.5, 1} is {0.9261, 0.8058, 0.7185} under $\boldsymbol{x_1}$ and {0.9395, 0.8174, 0.7274} under $\boldsymbol{x_2}$. It is evident the higher the Φ is, the more likely the consumer *i* will choose in the market. The reason behind this will be

[cx] The simulation was conducted with 1000, 3000, 5000, and 10000 iterations, separately. The results were found to converge as the number of iterations increases, with the one with more iterations exhibiting less random error. Since the goal of the analysis is to understand consumer behaviour in the long run (i.e., with large amount of purchasing experience), the present simulation, as well as all of the following ones, are run over 10000 iterations as a proxy of infinite learning.

explored later. Moreover, the slightly higher probability of choosing nothing in $\boldsymbol{x_2}$ versus $\boldsymbol{x_1}$ for each level of Φ stems from the fact that the mean value of inhibition of return in $\boldsymbol{x_1}$ is lower than in $\boldsymbol{x_2}$. This lower inhibition is because options in $\boldsymbol{x_1}$, on average, are preferred to options in $\boldsymbol{x_2}$ in terms of preferences$_U$. The results suggest that the difference between inhibition values under $\boldsymbol{x_1}$ and $\boldsymbol{x_2}$ may yield a difference in the probability of opting out.

The results also reflect the fact that the probability of not making any choice from the market is significantly affected by market inference (represented by the value of Φ) mainly through its influence on threshold values. Figure 11 (c) shows how threshold values are influenced by products' relative ranks and the false consensus effect and how threshold values of all market alternatives increase with the social influence parameter Φ. Recall that a market alternative is considered by the consumer *i* at the choice stage only when its inhibition is below the associated threshold value. The higher the threshold values, the more likely an alternative is to be considered at the choice stage, holding the value of inhibition fixed. Therefore, the fact that the threshold value is highest for the middle option and when $\Phi = 1$ means that an item can be more inhibited but still considered because its amount of inhibition is more likely to be less than the threshold value. Moreover, Figure 11 (c) indicates that as the tendency to adopt the matching strategy gets higher (i.e., as Φ increases), consumers are more likely to allow middle options into the consideration set[cxi], giving middle options an extra advantage over other options. This bonus is reflected in the fact that middle options are more likely to be chosen when Φ is high versus low[cxii], as revealed by Figure 11 (d), (e), and (f).

[cxi] By the definition of the matching strategy, consumers choose an option whose rank in the product line corresponds to consumers' perceived rank in the population. Combined with the false consensus effect, consumers are assumed to believe the middle option is most desirable, and therefore raise its threshold.

[cxii] Of course, the probability differences can also be attributed to choice differences in the exploration trials, where choice is determined in proportion to the norm-based preference with some random noise subject to (1- Φ). Yet, since inhibition of return of market offerings is quite large, choice is rarely made, limiting the social influence from exploration.

3.4.2.3 Market share for different values of Φ

As displayed in Figure 11 (d), the first option in each context ought to attract more demand than other market offerings since it is most intrinsically desirable and does not inherit inhibition of return from others. Yet, Figure 11 (e) and (f) show that if consumers are likely to use the rank order decision rule during decision making (i.e., $\Phi > 0$), demand for the first option in each context is spread to the second, third and fourth option in the set, but not to the fifth option. This reflects that the false consensus effect, which increases the thresholds of middle options more, allows the middle options to be included in the choice set despite their inhibition. Thus, the tendency for the first three options in the context to gain almost equally high choice share increases as Φ rises (as shown in Figure 11 (d), (e), and (f)). This change in market share of options as a result of a change in value of Φ implies the presence of the compromise effect when $\Phi > 0$, since middle options, as revealed in Figure 11 (d), are in effect not that attractive in terms of their preferences$_U$.

3.4.2.4 The strength of the compromise effect

The detection of the compromise effect leads to a question about how strong the effect is for each level of Φ. This subsection therefore assesses the strength of the compromise effect using each of the three methods described earlier, namely middle proportions, absolute-share changes, and relative-share changes. By comparing the absolute market share of the middle option in the present and the benchmark model, the estimated strength of the compromise effect in $\boldsymbol{x_1}$ and $\boldsymbol{x_2}$ respectively is 0.3462 and 0.2885 when Φ = 0.5 and 0.3893 and 0.3376 when Φ = 1, with no compromise effect for Φ = 0. The between-model and between-context differences indicates that the compromise effect reduces as the tendency to adhere to internal cues, i.e., the preference$_U$, increases as well as with the unfavourability of a context.

Moreover, Figure 11 (e) implies that x_4 (in $\boldsymbol{x_1}$) becomes chosen more often when its relative rank approaches the middle due to context change. Specifically, in the case where Φ = 1, the absolute market share of x_4 grows from 0.0004 to 0.3327 as its rank drops to the 3rd in $\boldsymbol{x_2}$. This is primarily attributable to an additional advantage from the false consensus effect, since this growth in the absolute market share is not found in x_2 and x_3 as context changes from $\boldsymbol{x_1}$ to $\boldsymbol{x_2}$, with their percentage increases being merely

6.91% and 12.55% when Φ = 1, much smaller than the corresponding benchmark results of 41.95% for both. Also, this growth in market share due to a contextual change does not appear for x_4 when $\Phi = 0$. Figure 11(d) shows that the market share of x_4 in $\boldsymbol{x_1}$ drops of 0.59% after it becomes the third option in $\boldsymbol{x_2}$ when $\Phi = 0$. By reference to the benchmark results, which suggest the natural percentage growth of x_4 to be 41.95% due to the removal of x_1, the impact of the compromise effect quantified via this measure (i.e., absolute-share changes) is 0.8322 and 0.8110 for Φ = 1 and 0.5, respectively. Apparently, the estimated magnitudes of the compromise effect are substantially higher than those obtained with the middle-proportions approach. This difference reflects the fact that in addition to being sensitive to the size of the positive bias towards the middle option, the current measure reflects a relatively low estimated market share x_4 generated in context 1 in the present model, in comparison to the benchmark.

Furthermore, the preferential treatment for middle options may trigger context-specific choice orderings. Figure 11 (e) suggests that x_2 and x_3 are markedly preferred to x_4 under $\boldsymbol{x_1}$, but that all these three options come to be close to indifference when the market context moves to $\boldsymbol{x_2}$. To be more specific, the relative choice frequencies of x_4 to x_2 and x_3 respectively are 0.11% and 0.12% in $\boldsymbol{x_1}$ and 49.92% and 49.94% in $\boldsymbol{x_2}$ when Φ = 1. This conspicuous change is also discerned when the degree of social influences is moderate (Φ = 0.5), with the relative shares of x_4 to x_2 and x_3 respectively being raised from 0.17% and 0.19% to 44.23% and 50.00%. In contrary, in the no-social-influence case, the change in context causes the relative choice frequency of x_4 to x_2 and x_3 to drop from 22.22% to 5.62% and from 39.08 to 26.92%. Overall, comparing Figure 11 (d), (e), and (f), it is evident that the weaker the social influence is, the more consistent the choice is.

Note that the closeness between $PU_i^m(x_3)$ and $PU_i^m(x_4)$ leads the estimated probability of choosing consistently between them to be just 50% in the benchmark model, slightly below the estimation of 55.04%, 50.00%, and 50.05% under the proposed model with Φ = 0, 0.5, and 1. More importantly, the probability of choosing x_3 over x_4 in $\boldsymbol{x_1}$ and x_4 over x_3 in $\boldsymbol{x_2}$ is 36.48%, 99.81%, 99.88%, 50.00 % respectively under Φ = 0, 0.5, and 1 and the benchmark. Such a high contribution (> 50%) detected under conditions of high social influence implies a strong systematic bias towards the third option in the context, in contrary to random errors appearing in the benchmark model. In short, while false-consensus consumers overall seem to be slightly more consistent than the rational agent with random errors, the high probability of reversing relative frequency in favour of the middle option reveals a strong compromise tendency.

In summary, the compromise effect may, to various degrees, be present when consumers are biased by perceived social norms and attempt to use an inaccurate market inference to make a decision. The present simulation shows that if learning in the market is independent of social influences, i.e., $\Phi = 0$, the choice pattern will resemble the preference$_U$-maximising one. By contrast, false-consensus consumers in the proposed model exhibit a strong propensity to choose the compromise option under both contexts. Therefore, despite having stable preferences$_U$ and considerable market experience, an uncertainty over preferences$_U$ and the false consensus effect may bias choice. Consequently, consumers may act as if preference orderings of the pair $\{x_2, x_4\}$ and of $\{x_3, x_4\}$ were not well-defined in isolation.

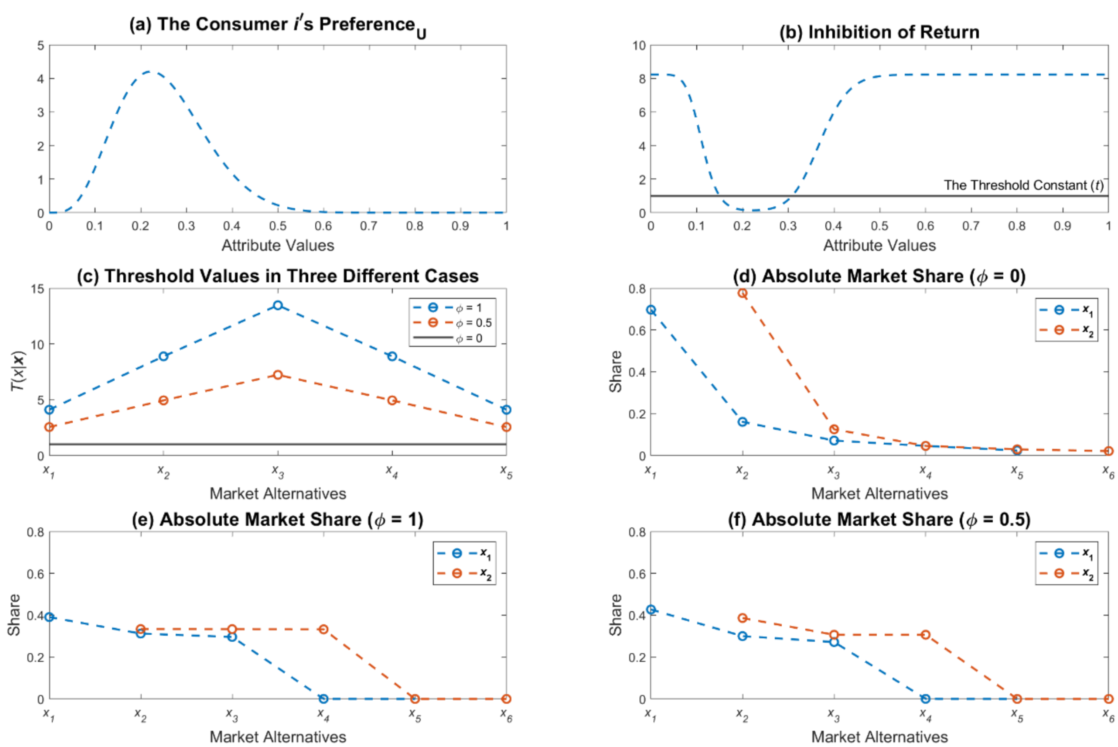

Figure 11. Simulation results on the demand side. (a) The representative consumer i's preference$_U$ over an attribute *m*; (b) The size of inhibition of return for every attribute value,

computed as $Inhibition_i(s^m) = 8.2253 \cdot e^{(-0.2312 \cdot (PU_i^m(s^m))^2)}$. Importantly, this is not a total amount of inhibition gained after a certain iteration. Rather, it indicates an incremental value of inhibition if a product is chosen. Therefore, it is computed without choice data. Moreover, the area between the inhibition function and the line of the threshold constant is the indifference area. Options with attribute values lying in this area will neither spread nor inherit any inhibition of return from others; (c) Values of the threshold $T(x|\boldsymbol{x})$ under full social influence, half social influence and no social influence cases, respectively. It is again, calculated before iteration started because computing threshold values only require a value of *t* and preferences$_N$. The bell-shape of the curves indicate that false-consensus consumers are more willing to give options that are ranked close to the middle a chance to be purchased again; (d), (e), (f) The average numbers of times that a product is purchased relative to other available market offerings throughout 10000 iterations, sorted by different levels of social influence Φ. Numbers presented in figures are averaged over total number of purchased in the market. Therefore, it is "market share", not "choice probability". The later denotes the number of times that a product is purchased divided by total purchasing periods (i.e., 10000 iterations).

3.4.2.5 The monopolist's optimal strategy and equilibrium contexts: Case 1

This subsection explores the monopolist's profit-maximising strategy, given consumer choice as described above, and the market equilibrium contexts in the case where products' marginal costs and prices are constant. Consider the case where the price *p* and the marginal cost *c* do not change with attribute value and their difference is strictly positive, i.e., $(p-c) > 0$. The profit function is thus written as $\Pi = (p-c) \cdot Q^D \cdot \sum_{j=1}^{n} Pr_j(\boldsymbol{x}; Q^D) - FC$, where $Q^D \in \mathbb{Z}$ refers to the total number of purchasing periods, $Pr_j(\boldsymbol{x}; Q^D) \in [0,1]$ indicates the probability of choosing x_j under context $\boldsymbol{x}$, given Q^D, and *FC* is denoted as an exogenously determined fixed cost. The probability, $Pr_j(\boldsymbol{x}; Q^D)$, is computed as the average quantity demanded of x_j across all experienced purchasing periods. Note that the probability of choosing an option is not directly linked to the market share of an option, in that the latter does not include data about choosing nothing in the market.

According to the present simulation results, over 10000 consumption periods, the total quantity demanded, as estimated under $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ respectively, is 2815 and 2726 when Φ = 1, 739 and 605 when Φ = 0, and 1942 and 1826 when Φ = 0.5. Given this, the estimated equilibrium context is $\boldsymbol{x}^* = \boldsymbol{x}_1$ for all three levels of social influence, if and only if the total market demand under $\boldsymbol{x}_1$ exceeds $\frac{FC}{(p-c)}$. Otherwise, the firm may not produce anything.

Moreover, the present results show that the firm has a higher monetary incentive to enter the market when consumers are subject to stronger social influence. Indeed, in either of the contexts, the monopolist enjoys extra profits from consumer bias. Mathematically, the actual size of the additional benefit depends on $(p-c)$ and Q^D. For instance, if the estimated probability, $Pr_j(\boldsymbol{x}; Q^D)$, is robust for all possible Q^D, the firm earns $(p-c) \cdot Q^D \cdot 0.2076$ units more in the equilibrium by selling $\boldsymbol{x}_1$ to consumers with Φ = 1, rather than to consumers who are immune to social influences. This is because the total quantity demanded of market alternatives is higher when Φ is higher due to that fact that the threshold value of each product is effectively increased by social influences (see Figure 11 (c)).

3.4.2.6 The monopolist's optimal strategy and equilibrium contexts: Case 2

The monopolist's optimal strategy and the equilibrium contexts in the case where products' marginal costs and prices depends on their attribute values are discussed in this subsection. Consider the alternative case where prices and marginal costs are both linearly associated with products' attribute levels. More specifically, suppose that $c_j = \delta_c \cdot s_j^m$ and $p_j = \delta_p \cdot s_j^m$, the profit function is instead written as $\Pi(\boldsymbol{x}) = \left(\delta_p - \delta_c\right) \cdot Q^D \cdot \sum_{j=1}^{n}\left[s_j^m \cdot Pr_j(\boldsymbol{x}; Q^D)\right] - FC$. The term $\sum_{j=1}^{n}\left[s_j^m \cdot Pr_j(\boldsymbol{x}; Q^D)\right]$ therefore determines the optimal context for the firm. Also, the amount of gains/ losses due to switching context is positively linked to Q^D and $\left(\delta_p - \delta_c\right)$.

With the current settings, the estimated results show that the summation of the product of each option's attribute value and its probability of being chosen, i.e., $\sum_{j=1}^{5}\left[s_j^m \cdot Pr_j(\boldsymbol{x}; Q^D)\right]$, under $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ respectively is 0.1381 and 0.1635 when Φ = 1, 0.0335 and 0.0326 when Φ = 0, and 0.0941 and 0.1081 when Φ = 0.5. In words, the monopolist can sell more "attribute values" under $\boldsymbol{x_2}$ than $\boldsymbol{x_1}$ in the case where consumers are affected by their perceived social norms[cxiii]. Accordingly, the firm has an incentive to set the context to be $\boldsymbol{x_2}$ when $\Phi = \{0.5, 1\}$ but $\boldsymbol{x_1}$ when $\Phi = 0$ when the aggregate attribute value sold over the whole periods exceeds $\frac{FC}{(\delta_p - \delta_c)}$. Otherwise, the firm is better off shutting down.

As shown above, if the firm chooses to enter the market it will earn more profits in equilibrium when consumer *i* tends to use market inferences to assist decision making. In this case, loss of quantity demanded due to offering a less attractive context is well compensated for by the additional gain from selling products with large attribute values. That is, the additional benefit from consumer bias motivates the firm to distort the informational content of its product line to take advantage of consumer bias.

3.4.2.7 Consumer welfare

Making the assumption that the values of preferences$_{\mathrm{U}}$ represent levels of happiness, this subsection explores average happiness gained from consuming in the market and potential welfare gain or loss as a result of a contextual change (or, the firm's manipulation of market context). A condition on the utility of an outside option for equalising consumer welfare under manipulated and non-manipulated contexts will also be investigated. Recall that it is assumed that consumers who do not choose in the market will instead choose an outside option. Note that the following analysis will not be true if consumers' preferences$_{\mathrm{U}}$ are uncertain and people are like the average person.

The simulation results indicate that estimated consumer welfare increases when the consumer is more resistant to the perceived social norms. More concretely, the average

---

[cxiii] Note that the social influence parameter Φ being strictly positive is not a necessary condition for market context $\boldsymbol{x_2}$ to be preferred. In fact, $\boldsymbol{x_2}$ turned out to be an optimal context in the "no social influence" case in some computer simulations when iterations = 1000, holding all other settings identical. As the magnitude of iterations grew, the finding that $\boldsymbol{x_1}$ is the profit-maximising context became more robust. In contrast, results obtained when $\Phi = \{0.5, 1\}$ are much more robust across many replicated simulations even when the number of iterations is just 1000.

happiness gained from consumption in the market under respective context is {0.8500, 0.1749} for Φ = 0, {0.5706, 0.0922} for Φ = 0.5, and {0.5314, 0.0810} for Φ = 1. Apparently, the average welfare level decreases with the magnitude of social influence parameter Φ[cxiv]. This implies that using market information incorrectly may reduce consumption-derived happiness.

Further, still considering welfare gain from consuming in the market, the strategy of shifting the context from $\boldsymbol{x}_1$ to $\boldsymbol{x}_2$ seemingly leaves the consumer *i* worse off. According to the present computational results, for Φ = {0, 0.5, 1}, the corresponding consumption welfare is decreased by 79.43%, 83.85%, and 84.76%. No surprisingly, percentage loss in welfare is larger in the biased-learning situation. Together with the fact that the profit-maximising context is $\boldsymbol{x}_2$ when Φ = 0.5 and 1, the findings implies that social-norm influences may further reduce consumers' welfare.

Yet, the existence of the outside option can possibly offset the welfare loss resulting from context change, in the sense that the consumer *i* is prone to opt out under $\boldsymbol{x}_2$. The condition on the utility of an outside option x_0 for equalising welfare under $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ is $PU_i(x_0)$ = {3.8983, 8.1015, 14.3277} for Φ = {0, 0.5, 1}. Recall that the maximum value of $PU_i^m(x)$ is merely 4.2034. This means that false-consensus consumers require very high utility from an outside option if they are to avoid excessive psychological loss due to social-norm induced biases as well as the firm's manipulation.

3.4.3 Simulation 2: Boundary conditions of the compromise effect

One of the key findings from Simulation 1 is that the degree of compromise effect falls as context moves away from the preference$_U$-maximising attribute value. The present subsection conducts a more detailed investigation of this phenomenon by extending Simulation 1 to the world where context changes in the way that attribute levels of all

[cxiv] This conclusion may be invalid if the utility of an outside option is considered. Since the probability of choosing in the market is positively related to Φ, the real welfare that consumers with Φ = 0 gain may be significantly reduced by a low level of utility of the outside option, while false-consensus consumers are far less affected.

current options are simultaneously raised by 0.01 units until the last option's attribute value equals to 1. That is, the computer simulation will start at $\boldsymbol{s}_{\boldsymbol{x_1}}^{\boldsymbol{m}}$ = {0.01, 0.11, 0.21, 0.31, 0.41} and end when it reaches $\boldsymbol{s}_{\boldsymbol{x_{60}}}^{\boldsymbol{m}}$ = {0.60, 0.70, 0.80, 0.90, 1.00}. For each context, there were 10000 purchasing periods. Note that unlike in Simulation 1, the notation x_j here will refer to the *j*th option in a context because there are more than two contexts simulated.

All other settings are identical to those in Simulation 1. Recall that parameter values were set as PU_i^m = Beta (5, 15), a decision-making threshold (*t*) = 1, an explore-exploit parameter (γ) = 0.5, a recency effect parameter (τ) = 0.9, the maximum inhibition periods (z) imposed to options with PU_i^m = 0 in a no social influence scenario = 20, and the parameter (g) that controls for the good-enough area, a.k.a. the indifference area, = 0.7. The rest of the subsection will compare the consumer i's choice under three different levels of social influence, Φ = {0, 0.5, 1}, and the benchmark model.

3.4.3.1 The benchmark predictions in choice

To begin with, given the PU_i^m which peaks at $s^m \cong 0.2222$, the preference$_U$-maximising choice x* of the consumer *i* is

$$x^* = \begin{cases} x_3 \text{ if } \boldsymbol{x} = \{\boldsymbol{x_1}, \dots, \boldsymbol{x_7}\}, \\ x_2 \text{ if } \boldsymbol{x} = \{\boldsymbol{x_8}, \dots, \boldsymbol{x_{17}}\}, \\ x_1 \text{ otherwise.} \end{cases}$$

This prediction from rational choice theory is that there will be no choice reversal between any two options observed across contexts. Given the possibility of random errors occurring during decision making, the benchmark market share of each option is computed by the softmax function, as shown in Figure 12. While the benchmark choice distribution follows the prediction of rational choice theory, in which options that maximise the preference$_U$ gain most shares, the market shares of all options estimated by the softmax function converge when the market context moves to the upper extreme of the attribute space, where the preference$_U$ is almost flat. As shown in Figure 12 the benchmark market share of x_1 declines gradually when the context goes beyond $\boldsymbol{x_{26}}$, whereas a slight increase in shares of other four options is seen.

3.4.3.2 Market share for different values of Φ

The market shares of options estimated by the proposed model for different value of Φ are discussed in this subsection. Like the benchmark model, Figure 12 shows that the market shares estimated by the present model with no social influence (Φ = 0) strikingly resemble the rational choice theory's predictions, even when the context is unfavourable. This implies that learning in the market can steer cognitive-constrained consumers toward preference$_U$-maximising choices. However, this observed similarity may not apply in the case of real choice probability as this observation is purely based on market share among options, which does not take probability of choosing nothing into account. Moreover, although x_1 seemingly dominates the market when the context moves away from the peak of the preference$_U$ in the attribute space, the probability of choosing it is quite small[cxv] and hence the difference of choice probability among market offerings is trivial. The observation of small choice probability stems from the high likelihood of opting out under contexts with large attribute values owing to high inhibition of return, as revealed in Figure 12 (a).

Prior to $\boldsymbol{x_{25}}$, the "M" shape of the curve observed in Figure 12 (a) reflects the fact that the optimal option shifts from x_3 to x_2 and then to x_1[cxvi] as context changes to include higher attribute values. More precisely, the probability of choosing an outside option[cxvii], x_0, climbs as the inhibition of return of x_3 rises, followed by a drop as x_2, which moves towards the maximum PU_i^m, replaces x_3 as the optimal product. As context changes, x_2 moves cross the peak of PU_i^m, leading the probability of choosing outside of the market to bounce back. This recovery continues after the best available option switches to x_1,

---

[cxv] Over 10000 iterations, the estimated probability of choosing x_1 is less than 5% under contexts from $\boldsymbol{x_{42}}$ onward.

[cxvi] This phenomenon does not appear in another two cases mainly because the threshold for considering an option is enlarged by the norm-based preference when $\Phi > 0$. The fluctuation of minimum inhibition of return due to context change is too small to influence probability of choosing in the market.

[cxvii] Remember that it is assumed that if consumers do not choose in the market, they choose an outside option.

while falling again as x_1 gradually departs from the maximum PU_i^m. Afterwards, when context passes a critical point, a dramatic increase in the optimal market option's inhibition of return precipitates a surge of choosing outside of the market.

Furthermore, Figure 12 illustrates that in presence of social influence, i.e., $\Phi = \{0.5, 1\}$, choice in the repeated market may nonetheless deviate from optimality, with choice under $\Phi = 1$ showing greater divergence. A glance of Figure 12 (b), (c), and (d) reveals that false-consensus consumers are less likely to choose the optimal option than unbiased ($\Phi = 0$) and benchmark consumers do[cxviii]. More specifically, when the context lies between $\boldsymbol{x_1}$ and $\boldsymbol{x_7}$, where the middle option, x_3, is the optimal option, the market share of x_3 is partly taken by x_2 and x_4 in the case of $\Phi = \{0.5, 1\}$ since these two options also benefited from the false consensus effect to some degree and incur moderate inhibition of return. Yet, as context moves towards the upper end, the benefit added to x_4 can no longer offset the large inhibition value it generates. Even worse, x_4 may receive inhibition of return spread from x_3 and x_2 as they gradually leave the "good-enough" zone. Therefore, when context enters the area where PU_i^m is rather flat, only x_2 and x_3 are observed to have an asymmetrical advantage. In particular, this positive bias increases with Φ and is more obvious on x_3, which helps it gain almost equal share to x_1 and x_2.

[cxviii] Note that for $\boldsymbol{x_{40}}$ onwards, the market share of x_1 estimated by the proposed model exceeds the benchmark results, making the predictions of the proposed model closer to the optimal (fully rational) outcomes. This is because differences in the preference$_U$ between x_1 and other alternatives get smaller as context becomes larger, amplifying randomness in the benchmark model while benefiting x_1 in the proposed model through the spill-over effect. Meanwhile, x_2 and x_3, which receive inhibition of return from options with lower relative ranks, obtain almost equal shares to x_1 thanks to social influences that raise the threshold values.

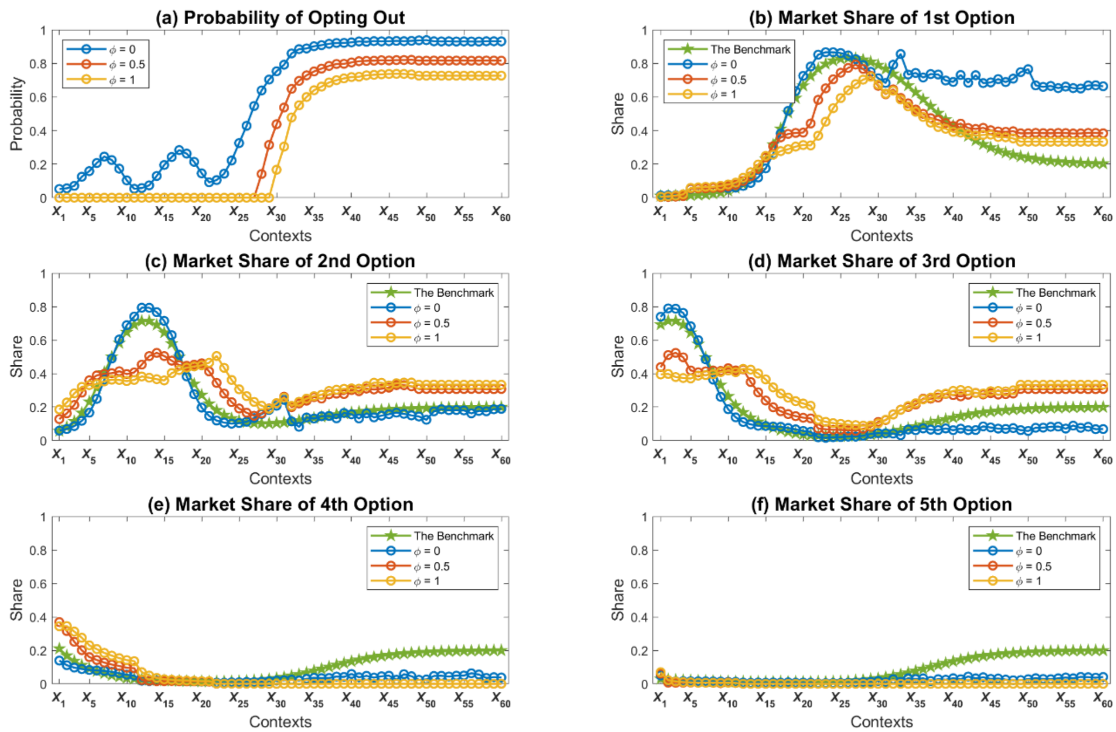

Figure 12. Consumer choice under each feasible context in the model. Plots (b), (c), (d), and (e) shows that the model with Φ = 0 behaves similarly to the benchmark model, whereas the models that include social influence by setting Φ = {0.5, 1} behave very differently.

3.4.3.3 The strength of the compromise effect

The extent to which choice is biased towards the middle option, x_3, is presented in Figure 13 and discussed below. Overall, the figure suggests that the degree of the compromise effect and its pattern vary with measures, with the middle proportion measure showing a downward trend whereas the rest indicate an upward shift. This difference highlights the effect of choosing different measurement approaches.

Starting with the most straightforward measure, the middle proportion approach, Figure 13 (a) illustrates that $\boldsymbol{x_8}$ is the watershed of the relative strength of the compromise effect. This is because prior to $\boldsymbol{x_8}$, the middle option, x_3, is the optimal choice, which makes consumers with Φ = 0 choose it more frequently than other options. For false-consensus consumers, choice of x_3 is dispersed to x_2 and x_4 since their corresponding preference$_U$ values are satisfied with the enlarged thresholds. Hence, the seemingly high concentration on x_3 simply reflects the fact that the model with $\Phi = 0$ is closer to the rational choice model in terms of its predictions. Contrarily, after $\boldsymbol{x_8}$, where x_3 is no longer the optimal option, the models with strong social influence show a preference towards x_3. This bias falls slightly when the context passes $\boldsymbol{x_{21}}$ owing to a sufficiently low preference$_U$ value of x_3. To sum up, Figure 13 (a) implies that the compromise effect may be mitigated as context becomes unfavourable.

Interestingly, a reverse pattern is observed in Figure 13 (b), in which the degree of the compromise effect increases when context goes beyond $\boldsymbol{x_{31}}$. This discrepancy mainly stems from the fact that the measure here is substantially influenced by the market share of x_4, rather than just depending in choices of x_3. As shown in Figure 13 (d) and (e), in the cases of $\Phi = 0.5$ and 1 the market shares of x_4 in the first ten contexts and those of x_3 in contexts between $\boldsymbol{x_{11}}$ and $\boldsymbol{x_{20}}$ are alike, yielding a small change in the market share of x_4 due to a change in the relative position. However, a slight increase in the market share of x_3 at $\boldsymbol{x_{30}}$, with that of x_4 remaining low from $\boldsymbol{x_{20}}$, triggers a noticeable change in the estimated absolute share. In particular, since the benchmark market shares of x_3 and x_4 gradually converge to each other, the estimated degree of the compromise effect for Φ = 0.5 and 1 becomes stronger as context moves towards the upper end, as presented in Figure 13 (b). In contrast, the degree of compromise effect estimated by the model with $\Phi = 0$ stays around zero[cxix] for almost all feasible contexts because the estimated patterns on x_3 and x_4 behave similarly to the benchmark results, as indicated in Figure 13 (d) and (e). In conclusion, the upward trend of the size of the compromise effect observed in the proposed model with $\Phi = 0.5$ and 1 shows that the absolute share of x_3 in a context $\boldsymbol{x}$ relative to that of x_4 in a context where x_3 in $\boldsymbol{x}$ is presented as the 4th option increases, whereas the difference between their preference$_U$ values declines. This type of favouring of the middle option is not captured in the first measure.

[cxix] The fluctuation in the degrees shown in Figure 13 (b) reflects the randomness of market shares of x_3 and x_4 estimated by the proposed model. Therefore, in the case of $\Phi = 0$, the degree of the compromise effect fluctuates more evidently when the absolute market share of x_4 gets smaller, given the fact that a smaller base value produces a larger percentage change.

Finally, the third measure focuses on the issue of choice consistency. It compares the relative market share between x_3 and x_4 in one context with that between the same pair of options in the context where they become x_2 and x_3 in the set. Figure 13 (c) shows that benchmark choice is highly consistent prior to $\boldsymbol{x_{25}}$. This is due to a high estimated probability of choosing x_3 over x_4 in the first ten contexts, which reduces as context moves up. This then leads the benchmark probability of choice consistency to approach to 50% as the relative share between x_3 and x_4 approximates 50%. Notice that the inconsistency in choice does not necessarily imply a systematic bias. As shown in Figure 13 (d), choosing x_3 over x_4 in one context and x_3 over x_2 in the context where the pair $\{x_3, x_4\}$ becomes $\{x_2, x_3\}$ due to a contextual change contributes the same amount to the overall inconsistency as the reverse case does. This means that benchmark consumers do not exhibit any bias towards the middle option. The choice inconsistency is purely due to random noise.

Moreover, regarding choice consistency, as illustrated in Figure 13 (c) and (d) a similar pattern is found in the case where $\Phi = 0$, while noise is observed due to randomness[cxx]. The sharp decline in consistency between context $\boldsymbol{x_{14}}$ to $\boldsymbol{x_{21}}$ is due to a dramatic decrease in the market share of x_3 in contexts with larger options than $\boldsymbol{x_7}$. As the relative share of x_3 and x_4 becomes stable across contexts, the probability of choice consistency plateaus at around 55%.

Contrary to results obtained shown in the above two models, the results in the cases where $\Phi = 0.5$ and 1 shows a systematic bias towards the middle option. Less abstractly, although the probability of choice consistency estimated under these two models is close to, and later slightly higher than, the benchmark estimate, Figure 13 (d) illustrates a strong tendency for false-consensus consumers to choose x_3 over x_4 in one context and x_3 over x_2 in the context where the pair $\{x_3, x_4\}$ becomes $\{x_2, x_3\}$ as a result of a change in contexts. Indeed, the seemingly more consistent choice observed in some contexts does not mean that false-consensus consumers suffer less bias than benchmark consumers do. Also, the

---

[cxx] The fluctuation in probability of choice consistency is expected to be reduced by increasing the number of iterations or consumers.

initially high inconsistency does not imply a great systematic bias towards x_3. Instead, the inconsistency in the beginning mainly arises from the fact that the relative share of x_3 over x_4 in the first ten contexts as well as that of the same pair in contexts between $\boldsymbol{x_{11}}$ and $\boldsymbol{x_{20}}$ is around 50%. As for contexts beyond $\boldsymbol{x_{35}}$, where choice appears to be as inconsistent as the benchmark but shows a strong bias, the main driver of inconsistency turns out to be the fact that false-consensus consumers are far more likely to choose x_3 over x_4 in one context, whereas choosing x_3 over x_2 in the context where $\{x_3, x_4\}$ becomes $\{x_2, x_3\}$ with a probability of approximately 50%. Therefore, the probability of choosing consistently is computed as 0.5×1, where 1 refers to the sum of relative share between x_3 and x_4 before the context change. Overall, despite seemingly low choice inconsistency across contexts, the degree of the compromise effect estimated with Φ = 0.5 and 1 grows with the context.

It is clear that the use of different measures can give different, and even opposite, patterns of the strength of the compromise effect. This implies that whether one takes the market share of x_2 and x_4 into account does matter. In addition, it is noteworthy that the whole computation is based on market share, not real choice probability. Therefore, a stronger compromise effect does not necessarily indicate that the middle option is more frequently chosen when attribute values in the context are large. Nonetheless, the results of the simulation imply that the compromise effect is highly likely to exist in equilibrium.

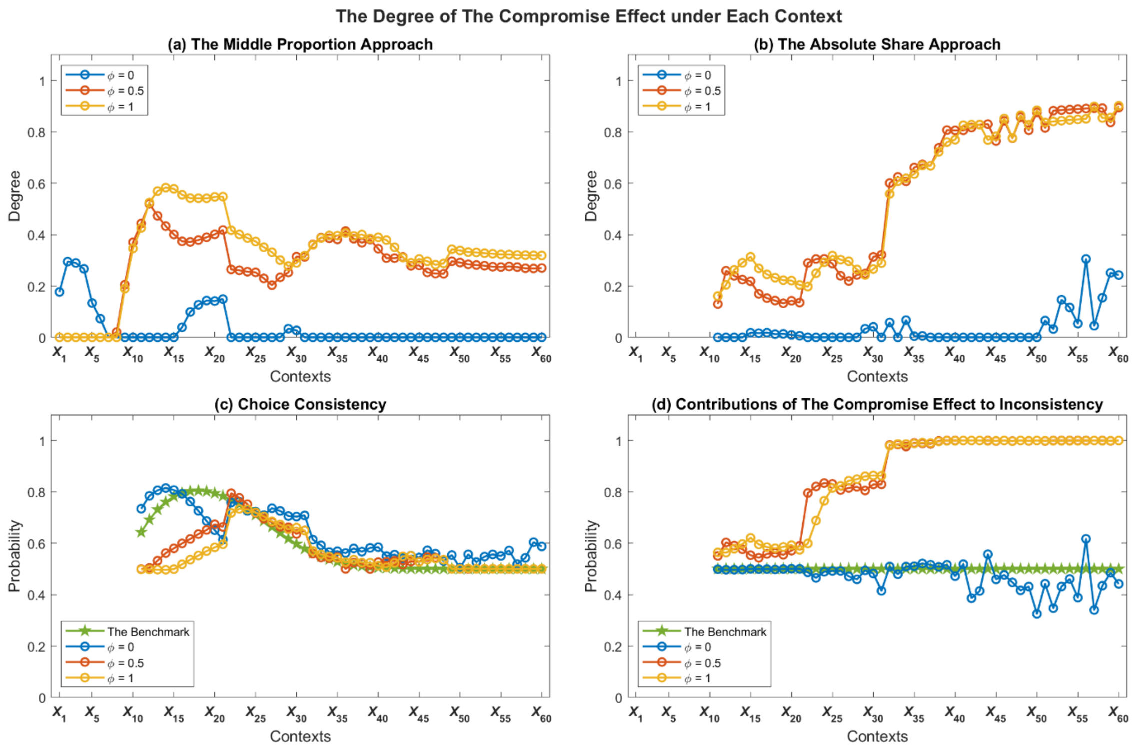

Figure 13. A graphical illustration of the degree of the compromise effect under each context. Plots (b), (c), and (d) start from $\mathbf{x}_{11}$ because they measure a change of absolute/ relative market shares from one context to the one located ten-contexts away. For example, results predicted under $\mathbf{x}_{11}$ indicate the difference between the market shares estimated under $\mathbf{x}_1$ and $\mathbf{x}_{11}$. Note that a pair of options $\{x_3, x_4\}$ in $\mathbf{x}_1$ is options $\{x_2, x_3\}$ in $\mathbf{x}_{11}$.

3.4.4 Simulation 3: The market equilibrium

This subsection extends the previous simulation with a focus on the firm's optimal response and consumer welfare in the equilibrium. It will explore the firm's strategy and equilibrium contexts in the market where only one type (characterised by Φ) of consumers exists. That is, the firm can design a context that is specifically offered to consumers with a certain level of Φ and consumers with other values of Φ cannot see that context because they are not in the same market. By separately investigating markets consisting of only one type of consumers, the effect of consumers' tendency to use (biased) market inferences on the firm's profit-maximising strategy and equilibrium contexts can be evaluated. The parameters and predicted consumer behaviour are same as in Simulation 2.

3.4.4.1 The monopolist's optimal strategy and equilibrium contexts: Case 1

This subsection discusses the monopolist's optimal responses to consumer demand and equilibrium contexts in the case where prices and marginal costs are fixed. It is also assumed that there is only one type of consumers in the market, whereby consumers with different levels of Φ are in different markets and therefore see different market contexts.

Assume that the price p and the marginal cost c is an arbitrary constant and $(p-c)>0$. Accordingly, the profit function is $\Pi=(p-c)\cdot Q^D\cdot\sum_{j=1}^{n}Pr_j(\boldsymbol{x};Q^D)-FC$, where $Q^D\in\mathbb{Z}$ represents the number of purchasing periods, $Pr_j(\boldsymbol{x};Q^D)\in[0,1]$ refers to the choice probability of an option x_j under context $\boldsymbol{x}$ given Q^D, and FC is an exogenously determined fixed cost. With fixed values of p, c, Q^D, and FC, the profit is maximised under contexts that generate maximum quantity demanded. Mathematically speaking, the optimal contexts also minimise the probability of choosing outside of the market. Therefore, a direct illustration of it is presented in Figure 12 (a).

Given the results estimated in Simulation 2, over 10000 consumption periods, a set of contexts that maximise total quantity demanded is $\boldsymbol{x}^* = \boldsymbol{x}_1$ when $\Phi = 0$, $\{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_{27}\}$ when $\Phi = 0.5$, and $\{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_{29}\}$ when $\Phi = 1$. This suggests the optimal contexts of market options that the firm should provide if and only if the total quantity demanded under $\boldsymbol{x}^*$ exceeds $\frac{FC}{(p-c)}$. Otherwise, the firm may be better off producing nothing. Note that since there is an overlap of profit-maximising contexts, there is no incentive for a firm to manipulate context information. In other words, the firm can enjoy earning the maximum possible profits by setting context $\boldsymbol{x}^* = \boldsymbol{x}_1$ for all types of consumers in the equilibrium.

Consistent with the findings of Simulation 1, the total quantity demand generated by the profit-maximising context, $\boldsymbol{x}_1$, is higher when consumers experience stronger social influences. This is captured by the fact that the probability of choosing x_0 under $\boldsymbol{x}_1$ is 0.0519 when $\Phi = 0$, but is zero when $\Phi = 0.5$ and 1. The greater quantity demanded gives the monopolist a stronger monetary incentive to enter the market as well as a higher possible profit it can earn in the market. Based on the present results, assuming the estimated choice probability, $Pr_j(\boldsymbol{x};Q^D)$, is constant for all possible Q^D, the firm can earn

$(p - c) \cdot Q^D \cdot 0.0519$ units more in the equilibrium by selling $\boldsymbol{x}^*$ to consumers with Φ = 0.5 and 1, rather than to consumers who are immune to social influences.

3.4.4.2 The monopolist's optimal strategy and equilibrium context: Case 2

The second way to address the profit-maximisation problem and equilibrium contexts is to assume that prices and marginal costs are linearly associated with products' attribute levels. Again, it is assumed that the market consists of only one type of consumer. Consider $c_j = \delta_c \cdot s_j^m$ and $p_j = \delta_p \cdot s_j^m$, the profit function is thus expressed as $\Pi(\boldsymbol{x}) = (\delta_p - \delta_c) \cdot Q^D \cdot \sum_{j=1}^{n} [s_j^m \cdot Pr_j(\boldsymbol{x}; Q^D)] - FC$. Here, it is the new term $\sum_{j=1}^{n} [s_j^m \cdot Pr_j(\boldsymbol{x}; Q^D)]$ that determines $\boldsymbol{x}^*$, instead of the total probability of choosing in the market alone. Moreover, the size of gains/ losses resulting from switching context is positively associated with Q^D and $(\delta_p - \delta_c)$.

As demonstrated in Figure 14, the summation of the product of each option's attribute value and its probability of being chosen, i.e., $\sum_{j=1}^{5} [s_j^m \cdot Pr_j(\boldsymbol{x}; Q^D)]$, is maximised under $\boldsymbol{x}^*$ = {$\boldsymbol{x}_{21}$, $\boldsymbol{x}_{27}$, $\boldsymbol{x}_{29}$} respectively for Φ = {0, 0.5, 1}. This indicates the optimal context for a firm to produce at when the aggregate attribute value sold exceeds $\frac{FC}{(\delta_p - \delta_c)}$. Importantly, $\boldsymbol{x}^*$ with Φ = 0 here is different from that in the first case in the sense that $\boldsymbol{x}^*$ is no longer the one that maximises $\sum_{j=1}^{5} [Pr_j(\boldsymbol{x}; Q^D)]$, while the small difference in $\sum_{j=1}^{5} [s_j^m \cdot Pr_j(\boldsymbol{x}; Q^D)]$ between context $\boldsymbol{x}_{21}$ and $\boldsymbol{x}_1$ implies that the effect of s_j^m on profits is limited. An increase in the attribute value provided easily offsets the loss in quantity demanded. However, when Φ = 0.5 and 1, the fact that $\boldsymbol{x}^*$ = {$\boldsymbol{x}_{27}$, $\boldsymbol{x}_{29}$} implies that the negative impact of the large increase in the probability of choosing outside of the market cannot be compensated for by the increased attribute values of available products. That being said, thanks to a high tolerance of somewhat unfavourable products, the monopolist is still able to sell more "attribute values" (and hence receive more profits) when consumers' choices are substantially affected by their perceived social norms. Consequently, in this case, the firm is strongly incentivised to distort the informational content of its product line to take advantage of consumer bias.

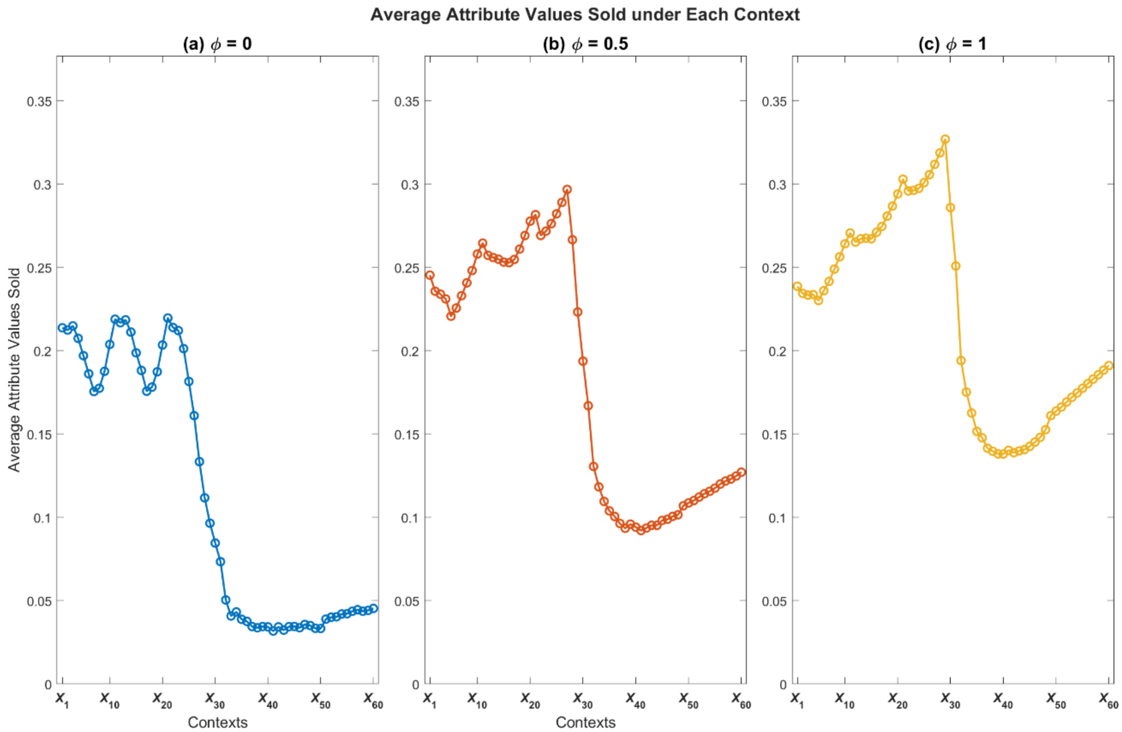

Figure 14. Attribute values sold under each market context, averaged over 10000 iterations. The observed reductions reflects an increase in the probability of choosing nothing. Comparing across the three graphs, the monopolist on average earns more when consumers are more likely to be influenced by their perceived social norms.

3.4.4.3 Consumer welfare

The manipulation of market context is of critical importance as it is inextricably intertwined with consumer welfare. Assuming that the values of preferences$_U$ represent the levels of happiness of consuming products, this subsection first investigates average happiness gained from choosing in the market and potential welfare changes as a result of the firm's manipulation of market context. It then explores a condition on utility of an outside option, x_0, for equalising consumer welfare under manipulated and non-manipulated contexts. Note that it is assumed that consumers who do not choose in the market will choose x_0.

I begin by examining the first case of the firm's profit-maximisation problem, where the monetary incentive to distort context is absent. As shown in Figure 15 (a), the socially optimal context is {$\boldsymbol{x_1}$, $\boldsymbol{x_{23}}$, $\boldsymbol{x_{24}}$} for Φ = {0, 0.5, 1}, where the average welfare gained from consuming in the market is {3.6454, 3.4919, 3.1096}, respectively. Accordingly, if we only consider welfare gained in the market, the estimate result implies that "rationally" using market information may reduce consumption happiness in the equilibrium if information is perceived with biases.

Furthermore, in the case where the monopolist is incentivised to manipulate the product line, the impact on consumer welfare is mixed. First, the firm intentionally sets context $\boldsymbol{x}^*$ = {$\boldsymbol{x_{21}}$, $\boldsymbol{x_{27}}$, $\boldsymbol{x_{29}}$} for consumers with Φ = {0, 0.5, 1}, to best respond to consumer bias. Again, considering only average welfare gained from consuming in the market, Figure 15 (a) reveals that switching the set of market offerings from $\boldsymbol{x_{21}}$ to the respective profit-maximising context decreases the average consumption welfare by 5.81% when Φ = 1, but increases it by 1.68% when Φ = 0.5. The growth in welfare in the case of Φ = 0.5 points to the possibility that biased consumers benefit from the firm's context manipulation. This is because although $\boldsymbol{x_{21}}$ is preferred to $\boldsymbol{x_{27}}$ in terms of the preference$_U$, i.e., $\sum_{j=1}^{5} PU_i^m(s_{x_j|\boldsymbol{x_{21}}}^m) > \sum_{j=1}^{5} PU_i^m(s_{x_j|\boldsymbol{x_{27}}}^m)$, choices made by false-consensus consumers do not meet this criterion well, with the extent of deviation being larger at $\boldsymbol{x_{21}}$ than at $\boldsymbol{x_{27}}$[cxxi] as shown in Figure 15 (a). However, the fact that consumers with the strongest social influences (Φ = 1) are worse off in the equilibrium implies that the unexpected "advantage" given by deviation in choice may fail to offset the negative effect of an unfavourable context on welfare. To conclude, the manipulation of context may not always reduce welfare gained from consuming in the market. Rather, its impact may depend on the degree of the bias.

[cxxi] This phenomenon can be observed in Figure 12 (b) and (c). In the context of $\boldsymbol{x_{21}}$, x_2 gains a high market share when Φ = 0.5 due to its low inhibition value and the preferential treatment imposed by (inaccurately) perceived social norms. Therefore, consumers with Φ = 0.5 exhibit a greater deviation from choosing the optimal option, x_1, in $\boldsymbol{x_{21}}$. Yet, when context moves to $\boldsymbol{x_{27}}$, the attribute value of x_2 is too undesirable, leaving it with high inhibition value. As a result, in $\boldsymbol{x_{27}}$, x_2 is unable to steal market share from x_1, whose inhibition value remains relatively low. The market share of the optimal option (x_1) thus increases, narrowing down the deviation.

Moreover, the evident welfare gap among different types of consumers observed in Figure 15 (a) implies that even if the monopolist does not manipulate the context, i.e., $\boldsymbol{x}^*$ = $\boldsymbol{x}_{21}$ for all Φ, false-consensus consumers will still suffer a welfare loss due to cognitive constraints and biased social sampling. This disparity purely results from the deviation from optimality.

Finally, the welfare obtained during the whole purchasing periods, including the situation of choosing outside of the market, is of equal importance to welfare gained from consuming inside the market. Figure 15 (b), (c), and (d) reveal how the utility of the outside option x_0 influences overall consumer welfare. The higher the utility of x_0 is, the more benefits consumers can receive from buying outside the market. In other words, the utility of x_0 affects consumers with Φ = 0 more strongly because false-consensus consumers are far less likely to choose outside of the market under each of all possible equilibrium contexts. Hence, consumers who are not influenced by social norms (Φ = 0) have the greatest advantage.

Note however that this argument cannot be made in the first case of the firm's profit-maximisation problem because, according to Figure 12 (a), the probability of choosing an outside option for consumers with Φ = 0 is zero under equilibrium context $\boldsymbol{x}^*$ = $\boldsymbol{x}_1$. Given that there is no incentive to manipulate context, it is expected that consumers who make choices purely based on their internal feelings, i.e., preferences$_U$, cannot benefit from (a high) utility of x_0. Moreover, since false-consensus consumers never choose outside of the market under their respective equilibrium contexts and there is no potential distortion of context, welfare of consumers with Φ = 0.5 or 1 are not influenced by the utility of x_0 either. In other words, the utility of x_0 does not affect the previously-noted analysis and conclusion regarding equilibrium consumer welfare in the first case, which suggests that the tendency to rationally make market inferences during decision-making may reduce consumption happiness in the equilibrium if information is perceived with biases.

With reference to the second case, the utility obtained from choosing x_0 is still not likely to compensate for the welfare loss resulting from the distorted context for consumers with Φ = 1. This is because the probability of choosing outside of the market in the equilibrium context, $\boldsymbol{x}_{29}$, is 0. Likewise, it cannot make $\boldsymbol{x}_{21}$, the equilibrium context for Φ = 0, be the welfare-maximising context of all feasible contexts for each of all types of consumers[cxxii] because when market context is at $\boldsymbol{x_{21}}$, only consumers with Φ = 0 choose x_0 sometimes, whereas other types of consumers always choose market offerings (see Figure 12 (a)). This limits the effect of utility of x_0 on welfare of consumers with Φ = 0.5 and 1. For consumers with Φ = 0, there are many contexts where consumers are more likely to opt out than under $\boldsymbol{x_{21}}$. Consequently, a large utility value of x_0 benefits those contexts more than $\boldsymbol{x_{21}}$, leaving consumers under the equilibrium context, $\boldsymbol{x_{21}}$, never enjoying the highest possible welfare level.

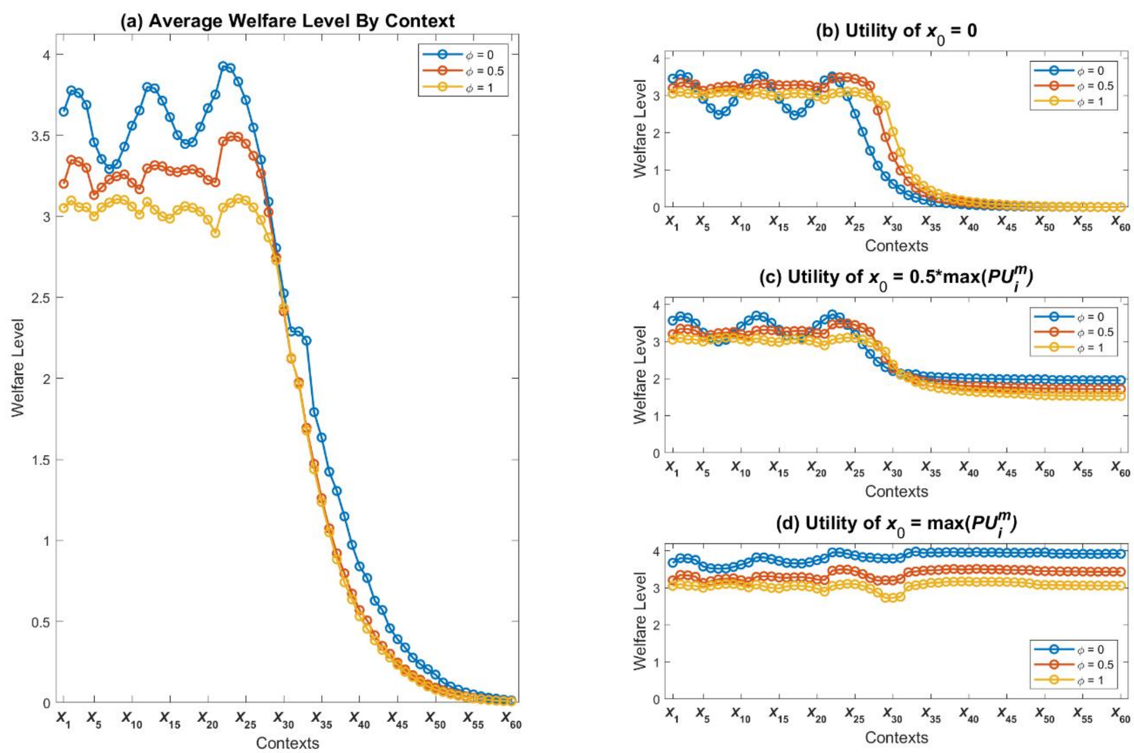

Figure 15. An illustration of consumer welfare under each context. (a) The average welfare level gained from choosing in the market is computed by summing up the product of market share of an option and its corresponding preference$_U$ value. Note that since the estimated welfare level here only captures those gained in the market, the slump in welfare level reflects a decline in the magnitudes of the choice's corresponding preference$_U$, not a decline in probability of choosing

---

[cxxii] Welfare levels for other types of consumers also matter as $\boldsymbol{x_{21}}$ represents the equilibrium context if the firm chooses not to take advantage of consumer bias.

in the market. The context that maximises average welfare from consumption for the cases where Φ = {0, 0.5, 1} is {$\boldsymbol{x}_{22}$, $\boldsymbol{x}_{23}$, $\boldsymbol{x}_{24}$}; (b), (c), and (d) Average consumer welfare with three possible levels of utility of the outside option. The difference between (a) and (b) reflects the fact that consumers with $\Phi = 0$ are at disadvantage when utility of x_0 is trivial. This is likely because the probability of choosing in the market is relatively low, even though choice itself is close to the optimal. Moreover, the main difference among these three figures appears when context lies in {$\boldsymbol{x}_{25}$, ..., $\boldsymbol{x}_{35}$}. The large change in welfare shown in (b) and (c), to some degree, corresponds to the one found in (a). As utility and choice probability of x_0 rises, the reduction gradually disappears.

3.4.5 Simulation 4: Multi-agent Scenario

To examine consumer behaviour and market equilibrium at an aggregate level, Simulation 4 created an artificial market with fifty unit-demand, heterogeneous consumers varying in the mean of their preferences$_{\mathrm{U}}$ distributions on the attribute m, i.e., the mean of Beta(α_i^m, β_i^m). Fixing the variance of all PU_i^m at 0.0089[cxxiii], the computer randomly selected the mean of each agent's PU_i^m from a symmetric distribution[cxxiv], Beta (5, 5). All values of the mean were redrawn if there existed any mean that caused $PU_i^m(x)$ to be infinity or made the derived $\alpha_i^m, \beta_i^m \leq 1$. In total, twenty sets of consumer preference$_{\mathrm{U}}$ were sampled and the simulation randomly chose the 9th set for analysis. The

---

[cxxiii] It is the variance of Beta (5, 15), the preference$_{\mathrm{U}}$ distribution used in the previous simulations.

[cxxiv] The authors note that in many real-world scenarios, the distribution of preferences$_{\mathrm{U}}$ in the population would likely be (positively) skewed; for example, this might apply to people's preferences for portion size, spiciness, cheese sharpness, etc. However, the overall finding and conclusion obtained in the case of skewed preference distribution is very close to the one explored in Simulation 2 and 3. To avoid repetition, this subsection concentrates on investigating the market where consumer preference is symmetrically distributed.

set of fifty consumers' preferences$_U$ curves used in this subsection is displayed in Figure 16.

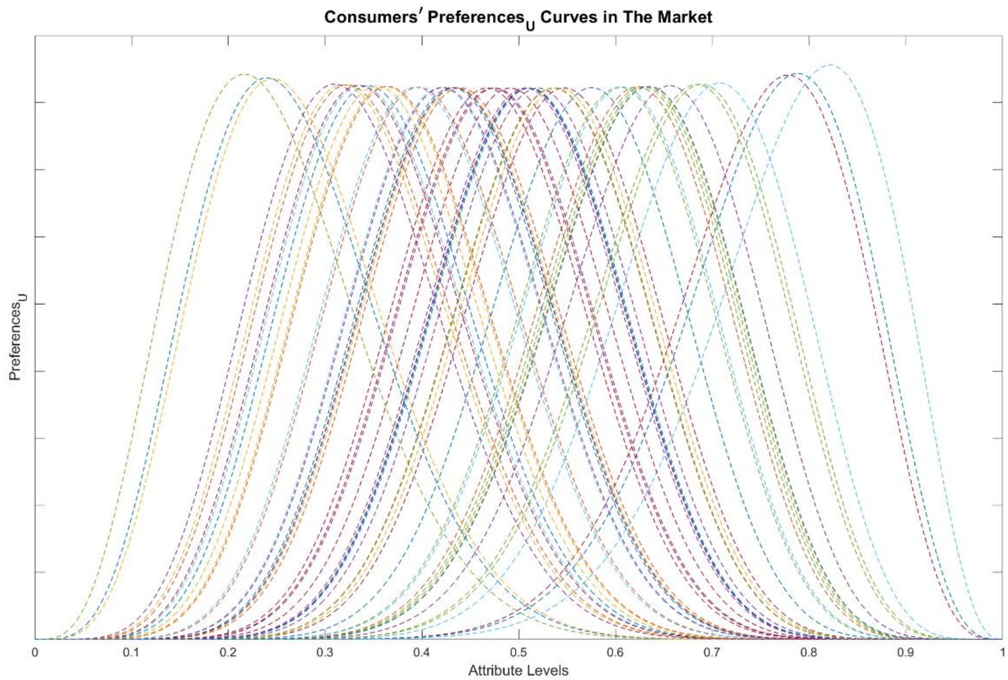

Figure 16. Consumer preferences$_U$ curves in the representative market. All distributions have variance of 0.0089 and different values of mean, randomly sampled from Beta (5, 5).

As before, the computer simulation was repeated 10000 times to approximate long-term consumer behaviour. The values of the parameters are identical to those used in Simulation 2. Contexts examined here are same as those in Simulation 2 and 3. More specifically, the simulation started from the context with $\boldsymbol{s}_{\boldsymbol{x}_1}^{\boldsymbol{m}}$ = {0.01, 0.11, 0.21, 0.31, 0.41} and ended at context with $\boldsymbol{s}_{\boldsymbol{x}_{60}}^{\boldsymbol{m}}$ = {0.60, 0.70, 0.80, 0.90, 1.00}. For each change in context, attribute values of all options are increased by 0.01 units. This subsection will compare consumer behaviour under three different levels of social influences, Φ = {0, 0.5, 1}, along with rational choice theory and the benchmark model. For convenience, this subsection sometimes uses the benchmark (optimal) model to refer to rational choice theory and the benchmark (noisy) model to denote the original benchmark model which is computed with a softmax function.

3.4.5.1 Market share estimated by each of all models

The market shares of options estimated by the rational choice theory, the benchmark (noise) model, and the proposed model with different values of Φ are explored in this subsection. How and why predictions from the latter two models deviate from the rational choice theory's estimation is also discussed.

First, Figure 17 shows the estimated market share of each option and the deviation from the prediction of rational choice theory. Despite a difference in the magnitude of market share, the choice pattern observed in each model is strikingly similar across models, whereby x_5 gains most of the share at the beginning but the demand gradually moves towards x_1 as attribute values in the context get larger. This pattern reflects the fact that when attribute values available are too small or too large overall, the market offerings with less extreme attribute values attract more demand since these less extreme products are associated with more favourable preferences$_U$ than other products in the menu. Notice that this difference in products' ability to attract demand under an extreme context is reduced by random noise. As illustrated earlier in Figure 12 in Simulation 2, when context is close to the extreme, market shares of all options will be nearly equal under the noisy benchmark model due to the small differences in products' corresponding preferences$_U$. Therefore, Figure 17 (a) shows that the benchmark (noisy) model has a U-shaped mean squared error curve.

It is noteworthy that the results estimated by the proposed model with zero bias are more consistent with the rational choice theory's predictions than those of any other models, as shown in Figure 17 (a). This indicates that unbiased learning in the market may help cognitive-constrained consumers make optimal choice more often. The change in the deviation observed in Figure 17 (a) simply reflects the fact that the optimal benchmark model is a deterministic model that is highly affected by switches of consumer choice. For example, as demonstrated in Simulation 2, a change of context from $\boldsymbol{x}_7$ to $\boldsymbol{x}_8$ would lead rational consumers to fully shift their choice from x_3 to x_2, while in all other models, market share of x_3 only declines by a few percentages. This characteristic makes the fluctuation in the market shares particularly noticeable under the optimal benchmark model, as shown in Figure 17.

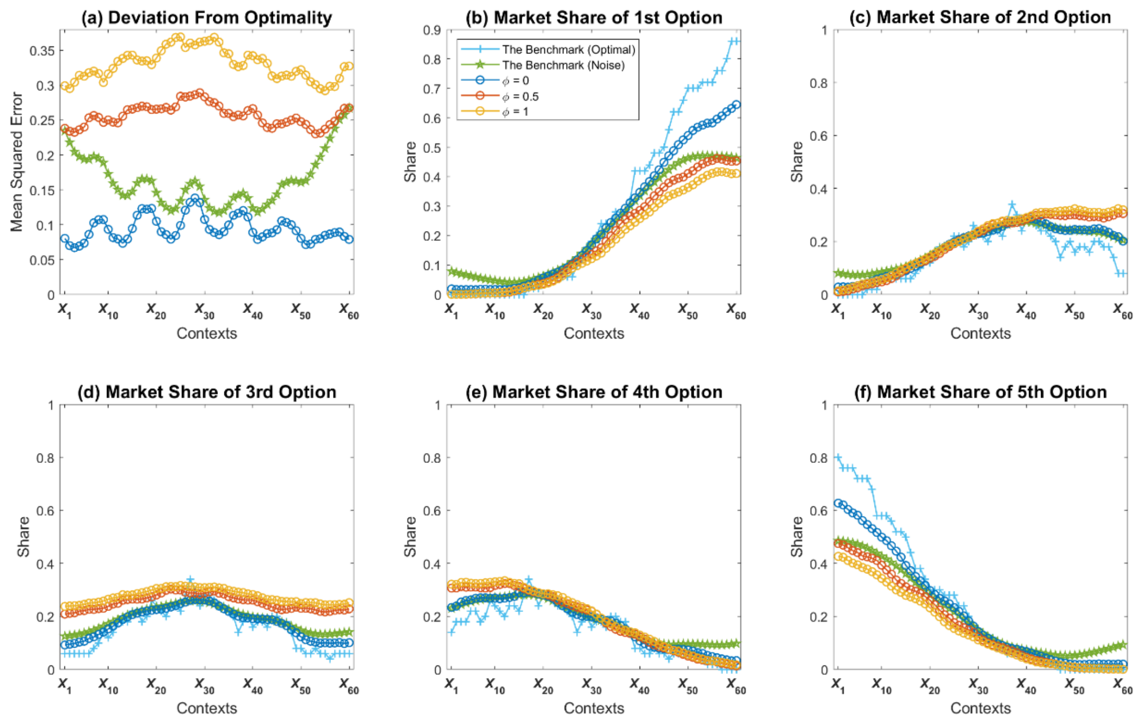

Figure 17. Market share of each option by context, averaged across consumers over 10000 iterations. The mean squared error was computed by subtracting market share of the best option estimated by the rational choice theory (the optimal benchmark model) from the one estimated by the model of interest, and then squaring it, followed by averaging out across consumers under each context.

However, this similarity in estimated market shares may not be generalised to the case of choice probability because the optimal probability of choosing the outside option x_0 is unknown. According to Figure 18, the probability of opting out, as well as its standard deviation, are high when context is extreme. This reduces the real choice probability of x_1 and x_5 under the small and large context respectively. If the optimal choice probability does not follow the same pattern, the proposed model will exhibit a greater deviation from optimality than the above observations do.

With regards to false-consensus consumers, i.e., Φ = {0.5, 1}, Figure 17 implies that biased perception of social norms may prompt consumers to exhibit extremeness aversion, in spite of their considerable consumption experience. More specially, the high tolerance of middle options underpins an abnormally high market share of x_3. As shown in Figure 17 (d), the effect is larger when market context is at the extremes. This is because most of the means of PU^m are located around the centre of $\boldsymbol{s^m}$, blurring the distinction between the most favourable option and the middle option under the extreme contexts. This enables x_3 to effectively take demand from the extreme, but more desirable, options, as shown in Simulation 1. To say it the other way around, when the context is set to be close to the middle of the attribute space, x_3 itself becomes optimal, giving social norms less room to exert influence. Interestingly, the deviation from optimality is higher when context is at a moderate level, as seen in Figure 17 (a). This counter-intuitive finding is rooted in the fact that almost all false-consensus consumers show a certain degree of deviation when context is located close to the middle of the feasible space, whereas a relatively smaller portion of consumers deviate when context is extreme[cxxv] due to a high inhibition of return for unfavourable products. The use of the mean squared error approach amplifies the impact of several small errors and therefore suggests that divergence is higher under the moderate contexts.

[cxxv] Correspondingly, the standard deviation of squared errors is larger when context is at the two extremes of the context space, compared to at the middle.

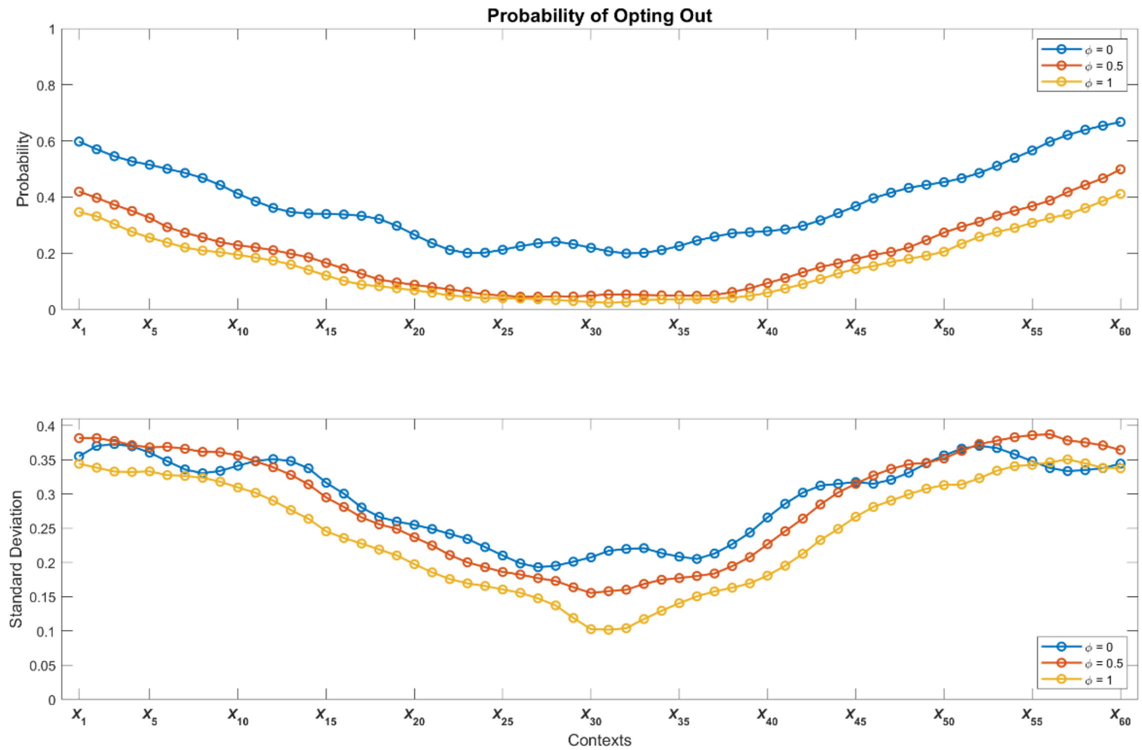

Figure 18. The probability of choosing the outside option and the standard deviation of that probability. The top figure suggests that false-consensus consumers are more likely to choose inside the market thanks to higher tolerance of unpleasant products, embodied as higher threshold values for all options. The fact that market quantity demanded is maximised at an intermediate context reflects the symmetric distribution of the mean of consumer preference. Less abstractly, when the mean of $\boldsymbol{s_x^m}$ gets closer to 0.5, more consumers will choose to buy in the market owing to a lower level of inhibition of return. By the same token, if the mean of market preferences$_U$ was drawn from a skewed distribution, the distribution of market quantity demanded would be skewed in the same direction as Figure 12 (a) shows. In addition, the bottom panel implies that the existence of consumers with more extreme tastes would make the standard deviation of probability of choosing x_0 under extreme context larger than the middle one. The, on average, higher level of standard deviation for the (Φ = 0) case implies a lower tolerance of such consumers for unfavourable products.

3.4.5.2 The strength of the compromise effect

The degree of the compromise effect is examined and presented in Figure 19. As noted in Simulation 2, the magnitude of the degree and its pattern vary with measurement approaches, with the absolute share approach showing a conspicuous upward shift (Figure 19 (b)) whereas the other two measures, namely middle proportions and relative-share changes, show a convex trend (Figure 19 (a) and (d)). Compared with the results in Figure 13, curves shown in Figure 19 look far smoother, because they capture the averaged behaviour of fifty heterogenous consumers.

3.4.5.3 Middle proportions

Beginning with the discussion of the middle proportion approach: Figure 19 (a) is consistent with Figure 17 (d), in that the compromise effect found in the case of positive social influences, i.e., Φ = 0.5 and 1, is strongest when context is at the extremes. Not surprisingly, this observation is different from the finding in the case of a single consumer with positively skewed preferences$_U$ i.e., PU_i^m = Beta (5, 15) (Simulation 2), as revealed in Figure 13 (a). This is because in the present simulation, the existence of false-consensus consumers with negatively skewed preferences$_U$ effectively increases the market share of x_3 when the context is at the upper extreme of the attribute space, making the estimated degree of the compromise effect bounce back in the end. Also, as implied in Figure 12 (d), the market share of x_3 may be at a low level when context is moderate, due to increased relative attractiveness of x_2. As a result, it is expected that consumers with extreme tastes may together drive down the market share of x_3, making the degree of the compromise effect weaker. Conversely, Figure 19 (a) suggests that the market of consumers with Φ = 0 exhibits no compromise effect for almost all contexts. This can be explained by the spill over of inhibition of return, which makes the extreme options more likely to be chosen under the proposed model with Φ = 0 [cxxvi], except in the situation where x_3 is the preference-maximising option. In sum, the results derived from the middle proportion approach state that in the market with false-consensus consumers, the aggregate compromise effect is most intense under extreme contexts.

[cxxvi] Meanwhile, random noise slightly increases the benchmark model's market share of x_3 when it is not the best option.

3.4.5.4 Absolute-share changes

However, if the compromise effect is measured via absolute-share changes, a different pattern of the strength of the compromise effect for Φ = {0.5, 1}, as shown in Figure 19 (b). As suggested by Figure 17 (d) and (e), the observed upward trend in Figure 19 (b) mainly stems from the fact that the share of x_3 in the proposed model with Φ = {0.5, 1} does not fall as much as it does in the noisy benchmark model (due to social influence), while the share of x_4 is roughly the same across all models when context lies between $\boldsymbol{x}_{15}$ and $\boldsymbol{x}_{45}$. For contexts from {$\boldsymbol{x}_1$, …, $\boldsymbol{x}_{10}$} to {$\boldsymbol{x}_{11}$, …, $\boldsymbol{x}_{20}$}, the change in the absolute share[cxxvii] as contexts include increasingly large attribute values is similar in magnitude in the proposed models and the benchmark model with random errors, weakening the compromise effect at the beginning. Furthermore, as expected, consumers with Φ = 0 do not exhibit any compromise behaviour for all feasible contexts, since the relative shares of x_3 and x_4 behave similarly to the benchmark results. The only difference occurs in the share of x_4 when context $\boldsymbol{x}$ goes beyond $\boldsymbol{x}_{50}$, which is not captured in the measure. To conclude, the increase in the degree of the compromise effect found in the proposed model with Φ = 0.5 and 1 indicates that the size of the bias towards x_3 is exaggerated as the preferential treatment imposed on x_4 is not effective enough to overcome its undesirability.

3.4.5.5 Relative-share changes

This subsection discusses results using the last measurement approach of the compromise effect, namely relative-share changes. This measure investigates choice inconsistency in a pair of options between two contexts and the proportion of the observed inconsistency results from the tendency to choose the middle option in the set. To recap, the measure of choice (in)consistency relies on the comparison between the relative market share between x_3 and x_4 in one context and that between the same pair in the context where {x_3, x_4} becomes {x_2, x_3}.

[cxxvii] Recall that the change in share here particularly refers to the difference in the absolute share between x_4 in one context and x_3 in the context where x_4 becomes the middle option, x_3, in the set. Based on the settings, x_4 and x_3 are in options with the same attribute value.

Figure 19 (c) suggests that the benchmark results are more consistent than those of any other model. The slightly concave shape of the benchmark curve indicates that choice is more random (= 50%) when context is around the two extremes. This is because, for most consumers, the preference$_U$ curve is rather flat for the extremely small and large attribute values. The horizontal line located at the probability of 50% in Figure 19 (d) further confirms that the choice inconsistency is due to random errors. In other words, benchmark consumers are not systematically biased towards the middle option.

Likewise, the proposed model with $\Phi = 0$ shows no selective preference for x_3, as illustrated in Figure 19 (d). The probability of choice inconsistency, however, appears to be higher than the benchmark estimation, as shown in Figure 19 (c). This is because, for all contexts, the relative share of x_3 to x_2 and x_3 to x_4, respectively, is lower in the proposed model than in the benchmark one (see Figure 17). That is, the high relative share of x_4 to x_3 in one context cannot be fully translated into the share of x_3 to x_2 in the new context where the pair $\{x_3, x_4\}$ becomes $\{x_2, x_3\}$, causing choices to be inconsistent as the context changes. In addition, the slight increase in choice consistency under contexts around the two extremes stems from the existence of consumers with skewed preferences$_U$, which allow them to choose more consistently in the desirable contexts consisting of options with extreme attribute values[cxxviii].

Equally importantly, although Figure 19 (c) shows that choice consistency estimated by the models with $\Phi = \{0.5, 1\}$ resembles the results when $\Phi = 0$, Figure 19 (d) suggests that the inconsistency under the former models is possibly a consequence of a tendency to choose x_3 over x_4 in one context and x_3 over x_2 in the context where $\{x_3, x_4\}$ becomes $\{x_2, x_3\}$. More specifically, observations under $\Phi = \{0.5, 1\}$ in Figure 19 (d) can be explained by the conclusion obtained in Simulation 2. The small increase in the size of the bias as context goes beyond $\boldsymbol{x_{53}}$ may be attributed to the existence of consumers with

---

[cxxviii] This phenomenon does not appear in the benchmark model possibly because the random noises make choice less optimal. As shown in Figure 13 (c), choice is initially more consistent in the ($\Phi = 0$) model than in the benchmark one.

positively skewed $preferences_U$[cxxix]. In similar vein, the slightly higher size of the bias in contexts between $\boldsymbol{x_{11}}$ and $\boldsymbol{x_{15}}$ is due to the behaviour of consumers with negatively skewed $preferences_U$. Consumers with more symmetrical $preference_U$ curves also, to some degree, contribute to the larger bias observed in the contexts consisting of options with extreme attribute values. On the whole, despite a moderate degree of choice inconsistency, the occurrence of choice reversal under the case of Φ = {0.5, 1} occurs mainly because of a systematic bias towards the middle option.

In summary, different measures produces different patterns of the intensity of the compromise effect. This underlines the importance of considering the market share of x_2 and x_4, in addition to x_3. Also, the results of the computer simulation imply that the compromise effect is highly likely to be present in the market equilibrium.

[cxxix] According to Simulation 2, when market context becomes unfavourable, consumers with positively skewed $preferences_U$ may more frequently choose x_3 over x_4 in one context, yet choosing x_3 over x_2 in the context where {x_3, x_4} becomes {x_2, x_3} with a probability of merely 50%.

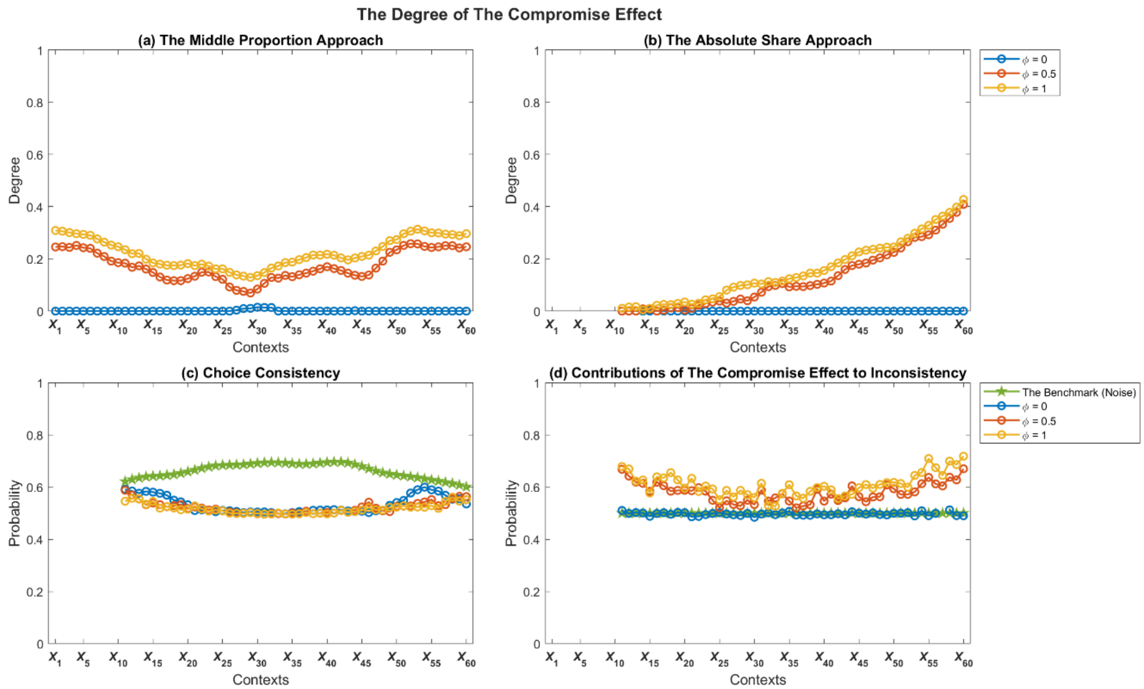

Figure 19. A graphical illustration of the degree of the compromise effect under each context. Note that for the first two measures (referring to panels (a) and (b)), instead of computing the degree of the effect for each consumer and then taking the mean, market shares were averaged out across consumers before computing the size of the compromise effect. This is because all resulting negative values would be treated as zero, leading to an underestimate of the effect of consumers who exhibit the "reverse" compromise behaviour. This would thus make the estimated degree larger than it ought to be. Moreover, as mentioned before, plots (b), (c), and (d) start from $\boldsymbol{x}_{11}$ since they measure a change of market shares of options with same attribute values. It is necessary to compare one context with one shifted by 10 increments because the context increases by 0.01 unit each time.

3.4.5.6 The monopolist's optimal strategy and equilibrium contexts: Case 1

To investigate consumer welfare in market equilibrium, the firm's profit-maximising contexts need to be identified in advance. As stated before, there are two possible cases of the profit-maximisation problem, in which the first case assumes a constant price p and marginal cost c. This subsection explores the firm's profit-maximising contexts and equilibrium contexts in the first case.

With a constant p and c, the aggregate profit function is written as $\Pi(\boldsymbol{x}) = (p - c) \cdot Q^D \cdot \sum_{i=1}^{50}[\sum_{j=1}^{5} Pr_{i,j}(\boldsymbol{x}; Q^D)] - FC$, where $Q^D \in \mathbb{Z}$ denotes the number of purchasing periods and $Pr_{i,j}(\boldsymbol{x}; Q^D) \in [0,1]$ refers to the choice probability of an option x_j under context $\boldsymbol{x}$ given Q^D for the consumer i. Based on the function, the profit is expected to be maximised at contexts that receive maximum total quantity demanded. Mathematically speaking, this is equivalent to minimising the mean probability of opting out. A graphical illustration of this is thus reflected in Figure 18.

As indicated by Figure 18, over Q^D = 10000 consumption periods, contexts that maximise market quantity demanded is $\boldsymbol{x}^*$ = {$\boldsymbol{x}_{32}$, $\boldsymbol{x}_{27}$, $\boldsymbol{x}_{31}$}[cxxx] for Φ = {0, 0.5, 1}. This suggests the contexts the firm should offer if and only if the total quantity demanded, $Q^D \cdot \sum_{i=1}^{50}[\sum_{j=1}^{5} Pr_{i,j}(\boldsymbol{x}^*; Q^D)]$, exceeds $\frac{FC}{(p-c)}$. Otherwise, the monopolist may be better offering nothing. Note that $\Pi(\boldsymbol{x}^*)$ is higher when consumers suffer stronger social influences, i.e., when they are endowed with a higher Φ. This is reflected by the fact that the probability of choosing x_0 is negatively associated with the values of Φ, as suggested in Figure 18. The larger quantity demanded, induced by higher Φ, gives the firm a stronger monetary incentive to enter the market and allows it to earn more profits in the equilibrium.

3.4.5.7 The monopolist's optimal strategy and equilibrium contexts: Case 2

This subsection discusses the firm's optimal responses to market demand and equilibrium contexts in the case where prices and marginal costs are no longer fixed. More specifically, the second case of the profit-maximisation problem assumes that prices and marginal costs are linearly associated with products' attribute values, i.e., $c_j = \delta_c \cdot s_j^m$

[cxxx] The fact of no overlap among profit-maximising contexts provokes a suspicion of context manipulation. However, there seems no specific direction of manipulation. As the set of the randomly sampled market preferences$_U$ change, the suspected "relationship" between profit-maximising contexts and the level of Φ changes. For example, under the 17th set, $\boldsymbol{x}^*$ = {25, 30, 32} for Φ = {0, 0.5, 1}, whereas $\boldsymbol{x}^*$ = {29, 30, 28} under the 4th set. Based on these results, there is insufficient evidence to argue that the firm may distort the context based on consumer bias, Φ, in this case.

and $p_j = \delta_p \cdot s_j^m$. The profit function is expressed as $\Pi(\boldsymbol{x}) = \left(\delta_p - \delta_c\right) \cdot Q^D \cdot \sum_{i=1}^{50}\{\sum_{j=1}^{n}\left[s_j^m \cdot Pr_{i,j}(\boldsymbol{x}; Q^D)\right]\} - FC$. As shown in Figure 20, profits are maximised at $\boldsymbol{x}^*$ = {$\boldsymbol{x}_{33}$, $\boldsymbol{x}_{38}$, $\boldsymbol{x}_{39}$} respectively for Φ = {0, 0.5, 1}. The profit-maximising firm will produce at these levels if the aggregate attribute value sold in the market with fifty consumers exceeds $\frac{FC}{(\delta_p - \delta_c)}$.

Apparently, the profit-maximising context in this case is located closer to the upper extreme of the set of feasible contexts than in the previous case for all values of Φ. This is because the profit is now determined both by available attribute values and by probability of choosing in the market. According to Figure 18, when Φ = {0.5, 1}, the probability of opting out is almost equally low at context $\boldsymbol{x}$ = {$\boldsymbol{x}_{22}$, …, $\boldsymbol{x}_{39}$}, making market attribute values more critical in generating profits. An increase in the attribute values compensates for the very small loss in quantity demanded. This strongly incentivises the firm to raise the overall attribute values of market offerings. Contrarily, when Φ = 0, the great fluctuation of probability of choosing an outside option, x_0, limits the relative effect of s_j^m on profits. This implies that the sharp decline in the probability of opting out cannot be offset by the increased attribute values of available products. Furthermore, as illustrated in Figure 20, the firm earns more profits under all feasible contexts when consumers are substantially influenced by their (inaccurately) perceived social norms. To conclude, in this case, there is a strong monetary incentive to distort the informational content of a product line to take advantage of consumer bias.

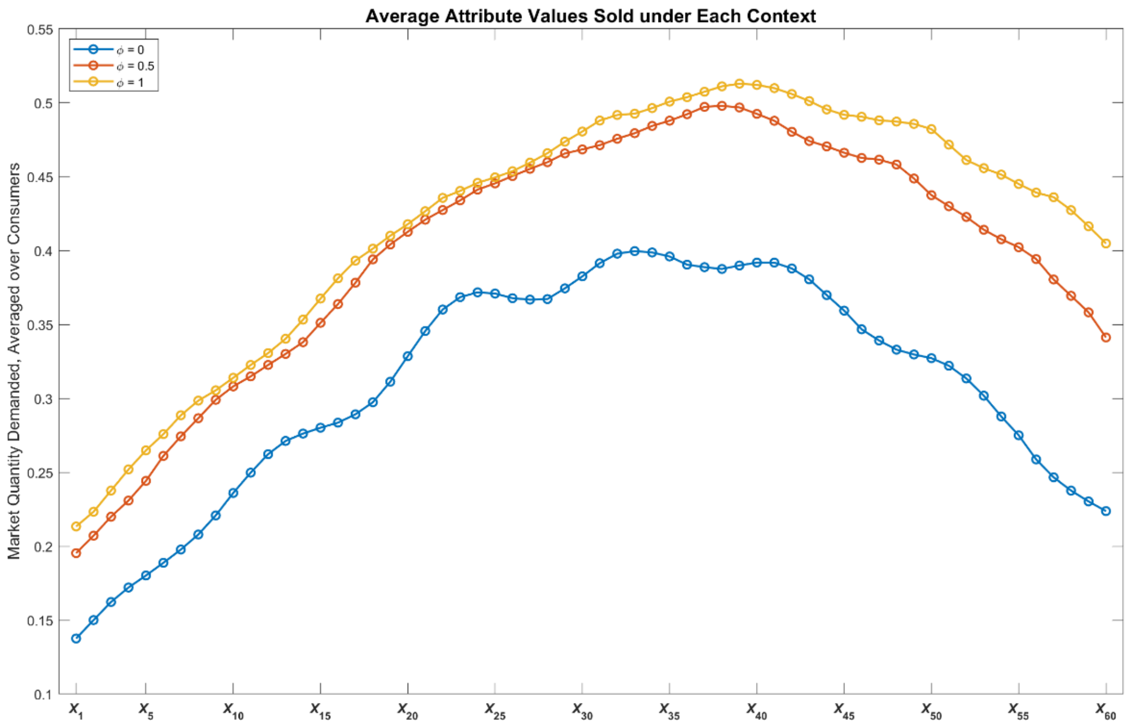

Figure 20. Attribute values sold by context, averaged across 10000 iterations and fifty sampled consumers. Comparing across the three lines, the firm earns more when consumers are more likely to be influenced by their perceived social norms. Moreover, the observed fluctuation, especially for the (Φ = 0) line, reflects the probability of opting out (see Figure 18). Specifically, the averaged attribute values sold in the market is inversely related to the probability of choosing the outside option. This perfect reflection indicates that the relative effect of attribute values sold on profits is limited when Φ = 0.

3.4.5.8 Consumer welfare

To examine whether the manipulation of contextual information will harm consumers, equilibrium consumer welfare is analysed in the following sections. This subsection first examines average happiness gained from consuming in the market and potential welfare change due to the firm's manipulation of market context. Since consuming an outside option x_0 also brings happiness to consumers, this subsection also explores a condition on utility of x_0 for equalising consumer welfare under distorted and non-distorted contexts.

First, in the first case of the profit-maximisation problem where price and marginal cost are constant, the average welfare gained from consuming in the market is {3.4639, 3.1145, 2.9191} under $\boldsymbol{x}^{*}$ = {$\boldsymbol{x}_{32}$, $\boldsymbol{x}_{27}$, $\boldsymbol{x}_{31}$} for Φ = {0, 0.5, 1}, respectively. Interestingly, although equilibrium context is not the same for different types of consumers, false-consensus consumers seemingly are, on average, not hurt from this kind of "manipulation". According to Figure 21 (a), for consumers who are subject to social influence, the averaged welfare level obtained from consuming in the market is relatively constant when context is between $\boldsymbol{x}_{27}$ to $\boldsymbol{x}_{36}$. The loss in consumption welfare here mainly stems from deviation from optimality owing to limited access to the magnitude of PU^{m} and social influences. Consistent with Figure 17 (a), deviation and thus disparity in welfare is greater when context is closer to the middle.

However, in the second case of the profit-maximisation problem, where product's prices and marginal costs linearly depends on their attribute values, a negative impact on consumption welfare is found. As argued earlier, the firm in this case is strongly incentivised to distort the product line in such a way that a context with larger attribute values will be offered to consumers with greater bias, i.e., consumers with higher values of Φ. The estimated equilibrium context is $\boldsymbol{x}^{*}$ = {$\boldsymbol{x}_{33}$, $\boldsymbol{x}_{38}$, $\boldsymbol{x}_{39}$} for consumers with Φ = {0, 0.5, 1}. Again, considering welfare gained from consuming in the market only, Figure 13(a) suggests that switching market context from $\boldsymbol{x}_{33}$ to the respective profit-maximising context decreases the average consumption welfare by 2.79% and 1.57% for Φ = 1 and 0.5. Unlike the results found in Simulation 3, an increase in deviation from optimality is no longer able to eliminate the negative effect of an unpleasant context on welfare. False-consensus consumers in this case are obviously harmed by the manipulation of context, even though they may have substantial purchasing experience.

Moreover, as consumers are allowed to choose outside of the market, the utility of the outside option, $PU_{i}(x_{0})$, plays a crucial role in equilibrium welfare analysis. Figure 21 (b), (c), and (d) show the effect of $PU_{i}(x_{0})$ on the overall welfare. Not surprisingly, the figure suggests that the magnitude of $PU_{i}(x_{0})$ affects consumers with Φ = 0 more since consumers with strictly positive Φ are less likely to choose x_{0} due to a higher tolerance of unfavourable products arising from social influence. Therefore, consumers who are uninfluenced by social norms (Φ = 0) receive greatest benefits from a growth in the value of $PU_i(x_0)$.

Finally, since consumers in the first case of the profit-maximisation problem do not suffer any welfare loss from context differentiation, the analysis on how $PU_i(x_0)$ improve equilibrium consumer welfare focuses on the second case. For consumers with Φ = 0.5, the required value of $PU(x_0)$ to equalise welfare[cxxxi] under the non-distorted context, $\boldsymbol{x}_{33}$, and the distorted equilibrium context, $\boldsymbol{x}_{38}$, is approximately 2.7153. In the Φ = 1 case, it is, however, not easily possible to reverse the welfare loss by simply having a high value of $PU(x_0)$ since the averaged probability of choosing x_0 under $\boldsymbol{x}_{33}$ exceeds that under $\boldsymbol{x}_{39}$, making $PU(x_0)$ unable to influence the outcome. This implies that the psychological loss due to the firm's manipulation is hard to avoid.

[cxxxi] Note that since for some consumers, $\boldsymbol{x}_{38}$ generates more welfare than $\boldsymbol{x}_{33}$ and/or the probability of choosing x_0 is greater under $\boldsymbol{x}_{38}$, it is possible to obtain a negatively infinite value of $PU_i(x_0)$ for those consumers. To avoid this situation, the computation took the market average of probability of choosing x_0 under each context of interest first. In other words, the analysis does not compute the value of $PU_i(x_0)$ for every consumer. The same method is used for all of the following analysis.

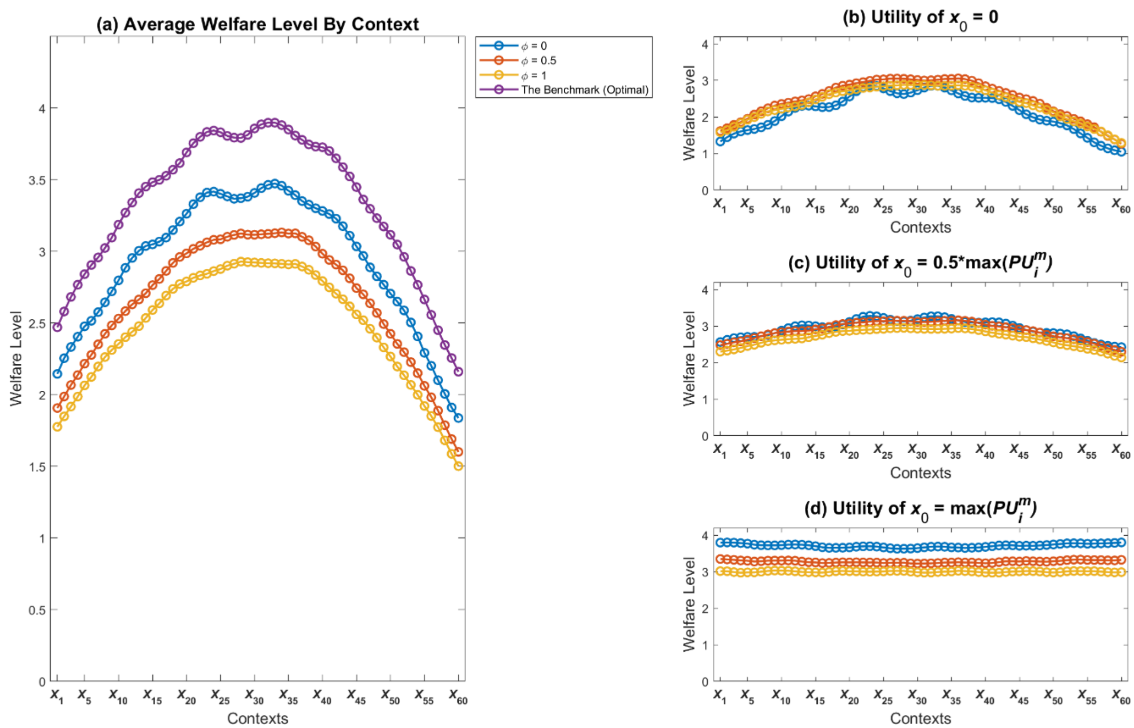

Figure 21. An illustration of consumer welfare under each context, averaged across fifty heterogenous consumers. (a) The average welfare level gained from choosing in the market, given the absolute market share of each choice and its corresponding preference$_U$ value. The context that maximises average welfare from consumption for the cases where Φ = {0, 0.5, 1} and the benchmark is {$\boldsymbol{x}_{33}$, $\boldsymbol{x}_{34}$, $\boldsymbol{x}_{28}$, $\boldsymbol{x}_{33}$}. The symmetrical shape of the curve implies that even consumers with extreme tastes enjoy a high level of consumption welfare when context is located around the middle of the attribute space. This can be seen, for example, in Figure 15, where consumers with PU^m of Beta (5, 15) gain the highest consumption-related welfare under contexts around $\boldsymbol{x}_{22}$ to $\boldsymbol{x}_{25}$. In addition, the magnitude of the difference in welfare level between the proposed model and the benchmark results implies the degree of deviation from optimality. (b), (c), and (d) Average consumer welfare with three possible levels of utility of the outside option. The difference between (a) and (b) reflects that consumers with Φ = 0 are at disadvantage when utility of x_0 is trivial. This can be explained by their relatively low probability of choosing in the market when Φ = 0 due to a relatively high threshold value for all options. However, the high probability of choosing outside may in turn improve welfare considerably as utility of the outside option raises to a certain level. Therefore, all of the curves become straight when $PU_i(x_0)$ is high, as shown in Plot (d).

3.5 Summary

This chapter presents an alternative model of context effects to account for consumer behaviour in the condition of repeated purchase. Several computer simulations were conducted to answer the following research questions. First, why do people with context-independent preferences$_U$ exhibit the compromise effect? Second, under what conditions will the compromise effect persist in market equilibrium, despite that fact that consumers have extensive market experience? More specifically, will consumers be more or less likely to exhibit the compromise behaviour when context is unfavourable? Further, is there any factor other than preferences$_U$ that affects consumer choice and can potentially lead to the compromise effect? Last, what is the firm's profit-maximising strategy in response to consumer bounded rationality and how does its strategic reaction affect consumers' welfares in equilibrium?

The second model is built upon the explore-exploit framework, in which false-consensus consumers at every purchase round will either explore, i.e., choose in proportion to the norm-based preference, or exploit, i.e., choose based on their inferred preferences. If choice is made from the market context, inhibition of return will be attached to the chosen option, with its amount being subject to the option's desirability. Products considered to be less favourable are assumed to inherit the same inhibition of return. Moreover, the inhibition size will accumulate but also decay over purchasing periods with a constant rate. The memory of choice history also decays. At each round, consumers will only consider options with accumulated inhibition sizes less than a threshold value. For false-consensus consumers, the threshold value of each option is norm-based. That is, the amount of increase of an option's threshold value depends on that option's rank in the set, with options located further away from the middle gaining less norm-incurred benefit. Therefore, false-consensus consumers are more tolerant of the middle options than to any other market offering.

This new model differs from the first model (Chapter 2) in three main aspects. Firstly, the learning and choice stage are no longer separated. Consumers in the second model never stop updating their inferred preferences and are sometimes willing to explore new options or retry options that have not been chosen for a long time, depending on their explore-exploit tendency. Choice is hence more flexible than in the first model, in which consumers are assumed to always choose the same option after pairwise learning. Secondly, social norms in the new model directly affect whether an option can be considered and hence whether to buy in the market. In the first model, these decisions are purely contingent on preferences$_U$. This change gives social norms extra power to influence purchasing decisions and hence affect the firm's optimal context setting. Thirdly, consumers in the first model are assumed to be unable to evaluate a product alone. In other words, owing to cognitive constraints, consumers only know their relative preference of a pair of products, but not the desirability of a product itself. In contrast, the second model adopts the idea of inhibition of return and assumes that consumers inhibit a chosen product and options that may be less favourable than the choice. This modification renders the strong assumption of pairwise comparisons no longer necessary, improving the plausibility of the model.

With the new model, the simulation results of both consumer and producer sides are consistent with expectations. Starting with the case where consumer(s)' preference$_U$ curve is positively skewed: The results suggest that market shares of products estimated by the proposed model with $\Phi = 0$, i.e., in absence of social influence, strikingly resemble the rational choice theory's predictions. This implies that learning in the market may lead cognitive-constrained consumers to choose optimally. Yet, it is also found that false-consensus consumers exhibit a propensity to choose the middle option when it is not the best option in the set. This is because when the middle option is optimal, its probability of being chosen may be transferred by perceived social norms to nearby options, which, in terms of the preference$_U$, are not much worse than the middle option. Other than this, as measured via middle-option choice proportions, the degree of the compromise effect is high for consumers who make inaccurate market inferences, i.e., $\Phi = 0.5$ and 1, with the intensity of the effect gradually declining as the whole context becomes more unfavourable.

In contrast, the measurement that assesses changes in absolute share reveals that the estimated strength of the compromise effect may increase with context's overall undesirability. This results from the fact that the absolute share of x_4 constantly falls with increased attribute values offered in market context, whereas the share of x_3 does not, due to social influence and less inherited inhibition of return. Results measured by changes in relative share demonstrate a similar pattern. While the difference in the probability of choice reversal between the benchmark and the proposed model seems to diminish as context gets closer to the upper limit, the contribution of the compromise effect to choice inconsistency, in the case of $\Phi = 0.5$ and 1, is always more than 50% and may even increase further as a function of changing context. This is because false-consensus consumers are predicted to choose x_3 much more frequently over x_4, but equally frequently to x_2 when all these options are not that desirable in terms of the preference$_U$. To conclude, although these measurement approaches suggest diverse patterns of effect sizes, they all suggest that the compromise effect may exist in market equilibrium where consumers can do repeated purchases. This implies that experience may not eliminate the compromise effect if preference learning involves bias.

Given consumer demand, the monopolist's profit-maximising strategy and equilibrium welfare were addressed in two scenarios. First, in the situation where price and marginal cost are not directly associated with attribute values, the simulation results show that there is no incentive to distort informational content of market context to exploit false-consensus consumers. In other words, the firm can enjoy maximum profits by offering the same context of options to all types of consumers. Consumer's welfare loss in equilibrium is therefore purely due to their estimation bias. However, if selling options with higher attribute values is more profitable (i.e., the second case), the firm's optimal strategy is to differentiate context provision, such that contexts with larger attribute values would be intentionally designed for consumers who suffer stronger social influences. This kind of manipulation reduces consumer welfare in the sense that false-consensus consumers are forced to face a less favourable context. Despite that, the simulation results suggest that consumers with $\Phi = 0.5$ are better off in the distorted context. This is probably because the undesirability of suboptimal options in the slightly worse context motivates biased consumers to choose the optimal one more frequently, while in the better context, the second- or even third-best option may be good enough to gain choice. The fact that the best option in the slightly worse context is better than suboptimal choice in the better context makes biased consumers better off in the slightly worse context. Notice that this conclusion largely depends on the parameters of the context such as the number of options, spacing between adjacent options' attribute values, etc. as well as the likelihood of utilising market information, Φ. As shown in Simulation 3, consumers with Φ = 1 are much worse off in the distorted context, as predicted.

In addition, regarding a different market scenario where there exist heterogeneous consumers in the market and the mean of their preference$_U$ curve is normally distributed. The results of the simulation implemented with fifty randomly sampled agents (Simulation 4) echo findings reported above. On the consumer side, people with Φ = 0, on average, deviate least from preference$_U$-maximising agents, in comparison to the benchmark (softmax) agents and people who are influenced by (inaccurately) perceived social norms. Moreover, these (Φ = 0) consumers, as expected, do not exhibit any compromise behaviour with respect to the three measurement paradigms. False-consensus agents, however, in aggregate show a certain compromise tendency in all hypothetical contexts, irrespective of measurement approaches. More specifically, in terms of the "middle proportions" metric, the degree of the compromise effect is strongest when market context is closest to the two boundaries of the attribute space. In contrast, the alternative measure using changes in absolute share suggests that the effect size increases with averaged attribute values set in menu. As before, this is likely due to the substantial decline in share of x_4. Finally, regarding choice inconsistency, it is observed that all types of consumers in the proposed model are more likely to reverse choice than the benchmark (softmax) agents are. Yet, when Φ = 0.5 and 1, the inconsistency is more likely to result from the propensity to choose the middle option before and after context change. For the benchmark and (Φ = 0) consumers, choice inconsistency is solely a consequence of random noise.

Furthermore, on the producer side, the firm's optimal responses to market demand in the heterogenous multiagent cases (Simulation 4) are also similar to the ones observed in the skewed preference case (i.e., the case of a single consumer presented in Simulation 3). First, if profits are not directly tied with the sold product's attribute value, no decisive evidence for intentional context manipulation is found. Although, in this heterogenous-agent scenario, the menu must be set differently for different types (with respect to Φ) of consumers to maximise profits, there is no clear direction of the observed context differentiation – it may be just a product of randomness. More importantly, all types of consumers seemingly are, on average, not made worse off by this differentiation. Again, welfare loss here results from deviation from optimality owing to limited access to the magnitude of PU^m and biased estimation on social norms. Second, the results show that the firm is incentivised to raise the attribute values of all market offerings in response to consumer bias, when selling products with higher attribute values is more profitable. Slightly inconsistent with the findings in the skewed preference case (Simulation 3), both types of false-consensus consumers now are, on average, worse off due to the manipulation of context, even though they have tremendous purchasing experiences.

To conclude, based on the simulation results, it is argued that preference learning in the market is likely to help cognitive-constrained consumers make optimal choices as long as no systematic bias is involved. However, people may treat the menu of market options as informative about the marketplace and attempt to use (inaccurate) market information to facilitate decision making. As shown in the simulations, a biased market inference, deriving from flawed estimates of one's own relative standing in the population distribution, may prompt consumers to exhibit compromise behaviour, despite consumption experience. This contextual influence is observed even in a very undesirable context, with the effect size contingent on the measurement metric. This highlights a methodological issue in measuring the compromise effect.

Finally, the propensity to choose the middle options gives the firm an opportunity to take advantage of biased consumers by distorting the product line. It is noteworthy that while parameter values in all computer simulations were set as reasonably as possible, the conclusion arrived at here is based on just one set of parameters. The estimated results may change with different assumed parameter values. Therefore, how each parameter affects the model's prediction needs to be explored in more details. Besides, as pointed out in the literature (e.g., Riefer, Prior, Blair, Pavey, & Love, 2017; Stojić, Schulz, Analytis, & Speekenbrink, 2020), the explore-exploit decision may depend on outcomes of past choices, instead of being purely stochastically determined.

Chapter 4: Experiment

This chapter reports empirical tests of the main assumptions of the models described in the previous two chapters. Full tests of the models' predictions are beyond the scope of this thesis; the aim here is simply to examine the plausibility of the basic assumptions. More specifically, the aim is to determine whether support can be found for the following assumptions of the model.

The false consensus effect. The false consensus effect is the tendency to believe that one's own attitudes and behaviours are more widespread amongst others than is really the case (e.g., Marks & Miller, 1987). The models presented in this thesis have assumed that an analogue of the false consensus effect applies to the specific case of *preferences*. The present study tests the idea that people tend to believe that they are near to average in terms of their preferences of a certain product attribute, even when they are not. This question is examined experimentally simply by asking people to state where they believe their preferences rank within the preferences of the population. To the extent that people's beliefs about their preferences are accurate, the distributions of these estimates should be uniform. If in contrast people are susceptible to a false consensus effect, the distributions will be unimodal with a peak near the median rank (.5).

Rank-order decision rule. A key assumption of the model is that consumers may have better awareness of how their tastes rank within a population than of what product attribute value they prefer. A further assumption is that consumers take markets of options to be informative about the distribution of preferences in the population (Wernerfelt, 1995), and a third claim, though not explicitly assumed in the models, is that consumers combine their beliefs about (a) their own relative rank and (b) the population distribution of choices over market options to choose a rank-appropriate option from the market. For example, if there are seven market options, and I believe that my preferences lie at the median of the population distribution, then I should choose the fourth market option if I also believe that population preferences and hence choices are symmetrically distributed over market options.

We test these assumptions by eliciting from each participant (a) their belief about where their own preferences rank within the population, as described above), (b) their beliefs about how population choices are distributed over the market options, and (c) their actual choice of product. If the basic model assumptions are correct, it should be possible to predict each individual's actual choices (i.e., c) from beliefs about their own rank (i.e., a) in combination with their beliefs about population preferences over market options (i.e., b). We attempt such prediction below.

This methodology does not, of course, enable us to identify causal relationships. To do so, one would need either to induce changes in people's beliefs about where they rank in the population (as Gershoff and Burson, 2011 do) or to change their beliefs about how population choices are distributed over market options (as many experimental manipulations of context do). It does however enable a basic "sanity check" of the model's assumptions. For example, suppose that a consumer believes her preferences lie at the median of the population distribution, but also believes that the population's distribution of choices over market options is positively (negatively) skewed. If the model's assumptions are correct, then she will choose an option that is lower (higher) than the middle option.

The experiment also enables evaluation of a number of additional, though less theoretically central, assumptions as follows.

Moderating role of product knowledge. The possibility that knowledge about a product category can mitigate the compromise effect in choice is also assessed in this chapter, although this hypothesis is neither assumed nor explicitly suggested in the results of the proposed models[cxxxii]. Note that results of past empirical studies (e.g., Sheng, Parker, & Nakamoto, 2005) suggest that knowledge and familiarity attenuate the compromise effect. Confirmation of such attenuation would be consistent with the intuition that "exploration-based" choices will be relatively less frequent than "exploitation-based choices" when experience of the product domain is higher.

[cxxxii] The two proposed models focus more about the role of preference learning or market experiences in diminishing the compromise effect, rather than knowledge about products.

Replication of compromise effect. An additional objective of this chapter is to confirm that the compromise tendency can be observed in a one-dimensional case, where an attribute-tradeoff is absent. This is done here simply by observing the distribution of participants' choices of options.

Choice and expressed preferences. In addition to the investigation of choice *per se*, subjective ratings of preferences$_E$ over products are elicited to verify whether the stated preferences are consistent with actual choice. It is anticipated that participants will choose the product that is assigned the highest preference$_E$.

To enable the above questions to be addressed, the experiment required participants to answer a series of questions regarding their preferences, levels of knowledge, estimation about their percentile standings in the population in terms of tastes, and purchase decisions over six distinct product categories, namely, cheese, curry dishes, steak, chocolate, pillows, and toothbrushes. A variety of product categories was used in the experiment to increase robustness of results since the measured compromise effect may vary with types of products (Neumann, Böckenholt, & Sinha, 2016).

4.1 Method

An online survey was conducted in March 2019. The survey was implemented using Qualtrics and took approximately fifteen to twenty minutes to complete. The experiment protocol was approved by the Humanities & Social Sciences Research Ethics Committee (HSSREC) at University of Warwick.

4.1.1 Participants

250 participants, all residents of the United States, were recruited from Amazon Mechanical Turk (MTurk) and electronically provided their informed consent to take part in the study. Each participant received $2.5 through their MTurk account for completion. Two were excluded from the sample since they may have submitted the survey twice[cxxxiii]. The final sample comprised 248 adult participants, with ages ranging from 20 to 72 years

---

[cxxxiii] This was detected by checking the IP addresses of participants.

(*M* = 38.19 years, *SD* = 11.39). 139 participants were male (56.05%), 108 were female (43.55%), and one did not disclose gender identity (0.40%). Participants varied in their ethnicity, with thirteen self-identifying as African, Caribbean, or Black American (5.24%), thirteen as Asian or Asian American (5.24%), ten as Mixed (4.03%), 207 as White (83.47%), and five choosing others or not to reveal (2.02%). Moreover, most of the sample (n = 152, 61.29%) were full-time employees and around half of included participants (n = 115, 46.37%) held a Bachelor's degree or above.

4.1.2 Questionnaires

The online survey required participants to complete several questionnaires, including a) demographic questions, b) prior ideal points, c) perceived positions in the population's taste distributions, d) knowledge about products, e) purchasing decisions, and f) expressed preferences. Except for demographics, all questionnaires were presented in random order[cxxxiv] over six product categories: cheese, curry, steak, chocolate, pillow, and toothbrush. The attributes used for these products were sharpness of cheese, spiciness of curry, degrees of steak "done-ness", proportions of cocoa in chocolate bars, softness of pillows, and softness of toothbrushes, respectively. Each product was defined in terms of one attribute only.

Prior ideal point was measured by asking participants to indicate their personal preference level of an attribute for a product category (e.g., spiciness of curry) on a 7-point scale. To prevent order effects, all seven questions were shown in random order. An example question about the sharpness of cheese is presented in Figure 22. Note that the scale represents the technologically feasible set of the attribute and, therefore, the prior ideal point is neither constructed on the space of available product nor relative-rank based. This measure hence is not an expressed preference as defined in previous chapters of the

[cxxxiv] Products were randomised but the types of questions were in a fixed order.

thesis, but something that will be used to (indirectly) test whether participants in general believe that the choice set reflects the technologically feasible set[cxxxv].

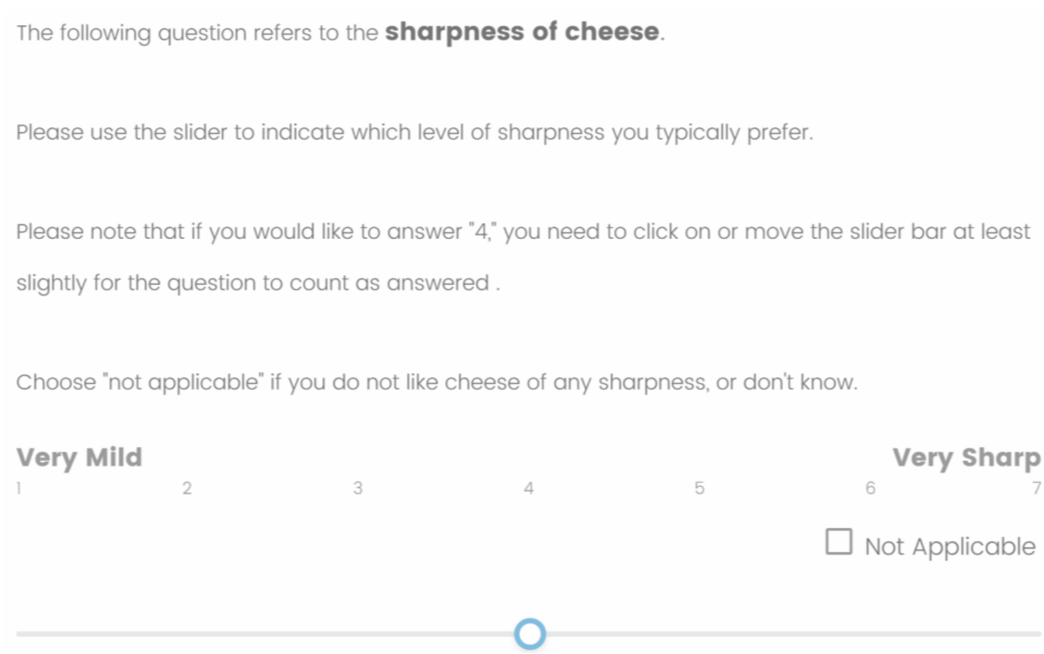

Figure 22. A screenshot of a question on sharpness of cheese in the prior ideal point questionnaire. The question asks participants to indicate which level of sharpness of cheese they prefer in general. Participants were allowed to choose “Not Applicable” to avoid the question, and those who chose this option did not see questions on the same product category in all following questionnaires.

Subjective relative positions in the population’s taste distribution were elicited by asking participants to indicate where they believed they were located relative to other people in the distribution of the population’s tastes for a product attribute (e.g. in preference for cheese sharpness). Questions were answered by clicking on a graphic of a

---

[cxxxv] This can be done by testing the correlation between participants’ prior ideal points and their choices in the market scenario, given the assumption that there is no other driving force of choices.

person in a line of twenty individuals (see Figure 23). Participants were told that these people were ordered in terms of their preference towards the product attribute, from liking a small attribute level to large. To improve the precision of the measure, the experiment followed the practice of Nisbett and Kunda (1985), whereby participants were explicitly instructed that the reference class they compared with was the U.S. population. Moreover, the use of a picture to present the question is different from the conventional methodology. For example, Burson (2007), Gershoff and Burson (2011)[cxxxvi] and Burson and Gershoff (2015), asked participants to estimate the exact number or percentage of people who are above or below them in their preferences.

[cxxxvi] In addition to point estimate, Gershoff and Burson (2011) also collected individual participants' perceived distribution of other people's performance on the quiz taking place in the beginning of the experiment by asking participants to allocate 100 points (i.e., people) to nine score bins. Unfortunately, this measure was not included in the present study as investigating the accuracy of the estimates of the population distributions on tastes in general is beyond the present scope.

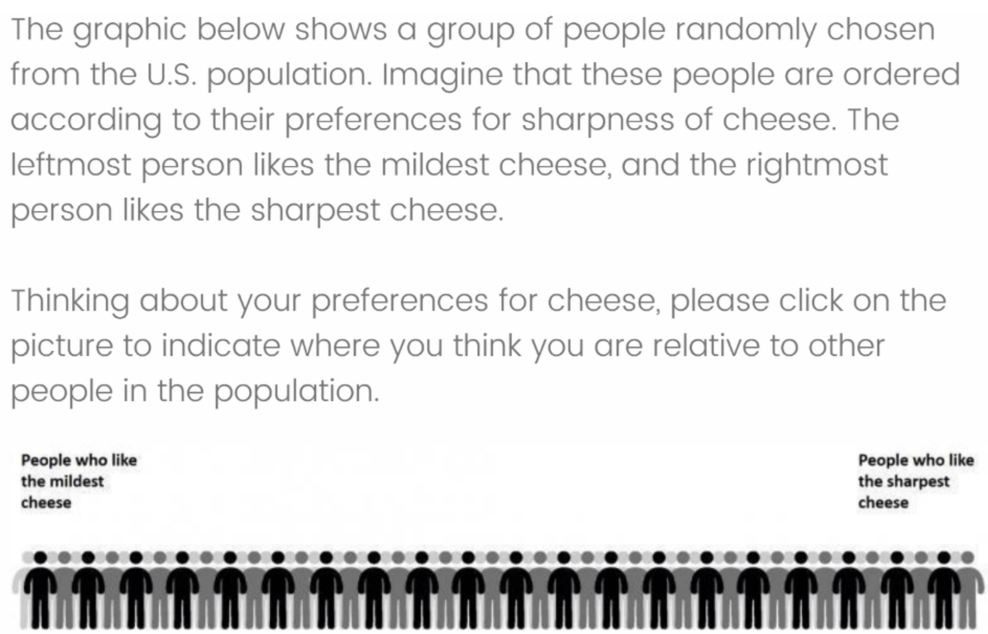

Figure 23. A screenshot of a question about estimate of one's own relative position in the population's distribution on the preference of cheese sharpness. In this questionnaire, participants were asked to choose their relative standings in the population by selecting one person from twenty U.S. citizens on the picture.

Knowledge about the product was assessed by asking participants how knowledgeable they consider themselves to be about each product, using a 7-point Likert scale ranging from 1 (not knowledgeable at all) to 7 (very knowledgeable). All questions were presented on the same page and in the same order ("cheese", "curry dishes", "steak", "chocolate", "pillows", and "toothbrushes"). Unlike with other measures, all recruited participants had to answer these questions. In other words, those who answered "Not Applicable" in the questions about the prior ideal points would still see these questions and could not avoid answering them.

This measure is a simplified version of Sheng, Parker, and Nakamoto's (2005) instrument, which combines three indicators to measure familiarity of products. In their paper, familiarity was elicited by asking about participants' familiarity about the product and knowledgeability about the two attributes presented in each product category.

Choice was elicited by asking participants to choose a product from a set of seven well-ordered (with respect to their attribute values) options, with an opportunity to pick nothing[cxxxvii]. This was the first of three successive questions in the last section of the survey, which measured participants' choices and preferences in a hypothetical market scenario. This section consisted of six question sets, each of which started with a brief introduction[cxxxviii] to encourage participants to imagine that they were shopping at a supermarket[cxxxix]. Underneath the introduction, participants were shown a picture of seven products ordered in terms of their attribute values, and were asked to choose one they would like to buy by clicking on the appropriate region of the picture (see Figure 24). Note that the only information provided to participants was the focal attribute of the first and the last product. Information such as brand names, prices, and specific attribute levels[cxl] were not available to participants.

[cxxxvii] According to Lichters, Sarstedt, and Vogt's (2015) review, experimental studies of the compromise effect usually employ a forced choice paradigm where participants cannot defer buying. However, the compromise effect may be more pronounced in the absence of the no-choice option possibly because choosing a compromise option and opting out are both ways to copy with difficulty of choice tasks (Dhar & Simonson, 2003). To avoid biased results, the present experiment allowed participants to make unforced decisions and participants who chose not to pick any product was eliminated from the relevant analysis.

[cxxxviii] The introduction was designed with the intention of providing participants with a concrete decision-making scenario following the suggestion of Alekseev, Charness, and Gneezy's (2017) that meaningful language, instead of abstract context, in experiments is useful in increasing participants' understanding of an environment and, thus, raise the quality of responses.

[cxxxix] For example, "Imagine that you are now shopping in a large, well-known supermarket, and are presented with a range of cheeses varying in sharpness. Suppose all cheese options are the same in all other ways (e.g., prices, shapes, etc.). Please answer the following three questions."

[cxl] Given that there is no attribute tradeoff in the choice task, disclosing the attribute values and product prices is not necessary in the experiment. This differentiates the present experiment from prior research on the compromise effect (e.g., Prelec, Wernerfelt, & Zettelmeyer, 1997; Sharpe, Staelin, & Huber, 2008; Sheng, Parker, & Nakamoto, 2005).

Figure 24. A screenshot of an example choice question. Participants were asked to choose one of the cheeses from an ordered set without being told exact attribute value of each product. The no-choice button on the right bottom reflects that participants were given a chance to defer choosing.

Expressed preferences were measured by asking participants to rate their preference of each available product presented in the previous question on a 5-point scale, namely "Not at all", "Natural", "Like a little", "Like a lot", and "Like very much". Note that there is no "Dislike" option since, mathematically, expressed preferences in our model are represented by a Beta distribution, which only produces positive density values.

Perceived taste distribution was obtained by asking participants to estimate market share of each available product by assigning twenty tokens (i.e., people) to them. An example of the question is shown in Figure 25. Participants provided estimates of frequency distributions of other consumers' choices. Each token in the task represented five people, and participants were asked to allocate tokens across seven levels of attribute value (e.g., seven levels of cheese sharpness).

The following question requires you to distribute twenty tokens to the columns below where each column corresponds to each cheese option that you faced in the previous question, with the leftmost column referring to the "Mildest" and the rightmost column referring to the "Sharpest."

Suppose there are a hundred people choosing to buy cheese in the supermarket. Please indicate the number of people that you think will buy each cheese option by pressing the "+" button below each column to distribute the tokens (1 token = 5 people).

For example, if you believe only five people will buy the mildest cheese, you should assign one token to the leftmost column. However, if you believe fifteen people will buy the mildest cheese, you should assign three tokens to the leftmost column. If you think no one will buy the mildest cheese, please do not assign any token to the corresponding column.

Once the task is completed, please click the "Submit the distribution" button to moved to the next page.

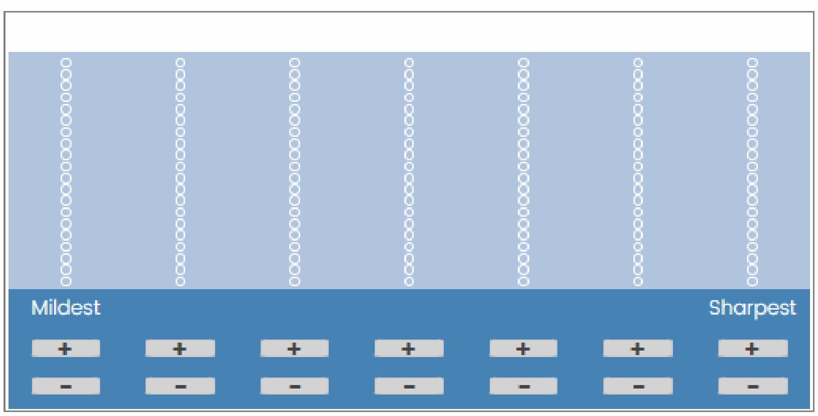

Figure 25. A screenshot of an example question about estimating the choice distribution of cheese. Participants were instructed to indicate their estimates of other consumers' choices by allocating tokens to each product. In total, there were twenty tokens, each of which represents five people. Products were presented in an order of their attribute values, with only the two extremes labelled.

4.1.3 Procedure

After completing a set of demographic questions, participants were asked to indicate their prior ideal points of product attributes, one for each product. Next, participants were presented with questions (in random order with respect to product types) about their estimates of their own relative standing in the population's taste distribution for each product category. The next section asked participants to indicate how knowledgeable they believed they were about each product and the last section elicited participants' choices in a hypothetical market scenario, where participants were asked to imagine they were shopping in a large, well-known supermarket. For each product category, participants were required to answer three successive questions, with one question per page. Firstly, they were asked to choose one product over a total of seven available products, ordered in terms of their attribute values. In the next survey page, participants indicated their preferences for every product and finally they were asked to assign market share of each product by allocating twenty tokens to products they saw in the first question.

The flow of the survey was designed to separate the prior ideal point questionnaire and the purchasing choice questionnaire to prevent rationalisation of choices by ratings. That is, the main purpose of placing two other questionnaires in between was to reduce the possibility that participants choose the product over the set just because its relative position is same as the point they indicate in the prior ideal point question.

4.2 Result and Discussion

Statistical analysis was performed using Matlab 2020b, developed by the MathWorks Inc. Throughout the analysis, level of statistical significance was set at $p < .05$. Importantly, since participants who chose "Not Applicable" for the prior ideal points question would not see follow-up questions on the same product category, numbers of data points used for analysis are different for each product category[cxli]. Specific numbers

[cxli] The only exception is questions about levels of product knowledge, which were presented to all recruited participants.

of observations are displayed in Table 2, along with descriptive statistics of the main variables of interest.

4.2.1 Descriptive statistics

Prior ideal points. As seen in Table 2, a majority of participants believed that their ideal points were located around the middle of the range of options, regardless of product category. The mode of either 3 or 5 (out of 7 points) implies that most participants felt that their tastes were moderate over the technologically feasible space. Only preference for cheese sharpness is close to the upper extreme, as indicated by a mode of 6. Moreover, distributions are mostly symmetrical; that is, the distributions of prior ideal points are rather bell-shaped[cxlii] as values of skewness are within the range of ±1[cxliii]. However, slightly large positive values of kurtosis (K) suggest that participants' aggregate stated preference have sharp peaks. Together, the two measures of a shape of a distribution do

---

[cxlii] The fact that the distributions are bell-shaped make prior ideal points more like the preferences$_N$ or the preferences$_I$, rather than the preferences$_U$. However, it is important to note that although questions that elicited prior ideal points were constructed on the technologically feasible set of the attribute and the reported ideal points were used as a proxy for the population's tastes in the later analysis, they are not the preferences$_U$, as the latter are not fully accessible and thus cannot be extracted before physically experiencing real products. Indeed, by definition, prior ideal points are closer to the preferences$_I$ or expressed preferences. Therefore, reported ideal points may similarly be affected by external forces such as social norms, which may be the reason why distributions of prior ideal points are symmetric as if such a distribution is a preference$_N$ itself.

[cxliii] Despite no general agreement having been reached in the literature yet, the present paper adopts Hair, Black, Babin, and Anderson's (2013) suggested cut-off values of skewness and kurtosis for normality assessment. In accordance to their criteria, the acceptable range of skewness is between ±1 and kurtosis is between ±2. It is also noted that there are other threshold values suggested in literature. For example, George and Mallery (2019) use ±2 for both measures, whereas Byrne (2013) argues that the range is better taken to be ±2 for skewness and ±7 for kurtosis.

not support normal univariate distributions of prior ideal points for five of the six product categories. Only the data on stated preferences of toothbrushes could be considered to be normal.

Knowledge. For most product categories, measures of central tendency[cxliv] indicate that most participants believed they were very knowledgeable about the product categories (see Table 2). The only exception is curry dishes where stated knowledgeability was relatively low. This result aligns well with the fact that many people did not indicate an ideal point for this product category. Moreover, the high K values, combined with a positive skewness, also implies that a majority of participants thought they knew the products well (again, with the exception of curry dishes). Last, data on knowledge is, on average, more concentrated and asymmetric than those on prior ideal points, as indicated by smaller standard deviations and larger absolute values of skewness and kurtosis for this question (see Table 2).

Table 2. Descriptive Statistics of the main measures for each product category.

	Present results								
Measures	*N*	*Min*	Median	*Max*	Mode	*M*	*SD*	Skewness	Kurtosis
Prior ideal point									
Cheese	244	1	5	7	6	4.68	1.55	-0.38	2.15
Curry dishes	217	1	5	7	5	4.17	1.83	-0.31	2.09
Steak	231	1	4	7	3	4.06	1.51	0.20	2.16
Chocolate	241	1	5	7	5	4.82	1.22	-0.42	2.90
Pillows	247	1	4	7	3	3.80	1.48	0.09	2.29
Toothbrushes	246	1	3	7	5	3.49	1.49	0.05	2.00

[cxliv] This refers to mean, median, and mode.

Perceived percentile in population									
Cheese	244	1	13	20	14	12.51	4.89	-0.40	2.15
Curry dishes	217	1	12	20	15	10.85	5.83	-0.27	1.81
Steak	231	1	11	20	11	10.88	5.06	0.11	1.95
Chocolate	241	1	14	20	11	13.16	4.05	-0.51	3.01
Pillows	247	1	10	20	11	9.91	4.71	0.09	2.08
Toothbrushes	246	1	8	20	7	8.91	4.75	0.26	2.12
Knowledge									
Cheese	248	1	5	7	5	4.72	1.51	-0.46	2.76
Curry dishes	248	1	2	7	1	2.92	1.75	0.54	2.13
Steak	248	1	5	7	6	4.92	1.61	-0.77	3.00
Chocolate	248	1	5	7	6	5.00	1.40	-0.63	2.99
Pillows	248	1	5	7	5	4.80	1.43	-0.24	2.49
Toothbrushes	248	1	5	7	5	4.98	1.34	-0.39	2.73

Notes: The table displays descriptive statistics of the main measures. Numbers of observations indicate that of 248 included participants, almost all provided their ideal points for each product category, except curry dishes received many fewer responses. Moreover, skewness and kurtosis are presented to characterise the shapes of the distributions.

4.2.2 The false consensus effect

This subsection tests the hypothesis that there will be a false consensus effect in preferences for a product attribute by analysing participants' estimated relative standings in the population with respect to their tastes about a product attribute. At first glance, Table 2 reveals a tendency for people to believe that they are close to average in the population distribution of tastes. For categories such as steak, chocolate, and pillow, the eleventh person (out of twenty representative people) was the most frequently chosen one.

This means that participants generally thought that they were at the centre of the population's distribution, consistent with a false consensus effect for preferences. This pattern is consistent but less pronounced for cheese, curry, and toothbrushes.

To illustrate participants' estimates about their relative tastes more clearly, we grouped responses into five bins, each of which represents 20% of the population. A set of histograms of the grouped data by product category is displayed in Figure 26. It can be seen that the middle three quintiles receive more choices than the extreme ones do. This reflects the fact that participants were less likely to state that their relative tastes were at the extremes of the population distributions. Moreover, the distributions of participants' subjective rank positions in the population for several product categories, such as curry dishes, steak, pillows, and toothbrushes, are found to be close to normal. On aggregate, the values of skewness for categories, namely cheese, curry dishes, steak, chocolate, pillows, and toothbrushes, are -0.44, -0.22, 0.05, -0.49, 0.08, and 0.28, and K for the same ordered categories are 2.09, 1.67, 1.91, 2.78, 1.94, and 1.97. The normality of the distributions for many product categories further supports the view that the majority of participants perceived themselves as close to the average in terms of their tastes.

In addition to the analysis of skewness and kurtosis, a more advanced normality assessment was conducted by using formal tests, including the Anderson-Darling test, the Jarque-Bera test, and the Lilliefors test[cxlv]. However, none of the distributions appeared

[cxlv] All these three tests are based on the null hypothesis that the data comes from a normal distribution, without requiring to set a value of population mean and variance. The analysis employed three distinct tests for normality assessments because each of them has its own merits. According to Yap and Sim's (2011) power comparisons, of the these chosen tests, the Jarque-Bera test is most powerful for symmetric long-tailed distributions and the Anderson–Darling test is best for asymmetric distributions. Although the Lilliefors test seems to have no advantage compared to another two, it still provides an additional reference to evaluate normality.

to follow a normal distribution[cxlvi]. Despite that, the overall results still provide partial evidence in support of the assumptions of the false consensus effect in human preferences.

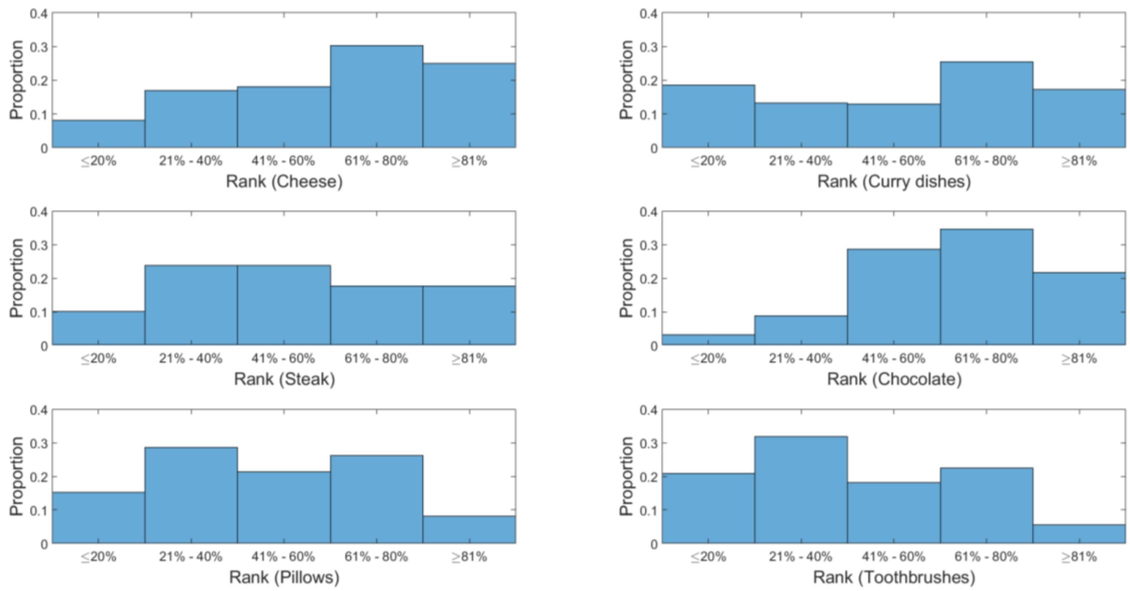

Figure 26. Participants' subjective relative ranks of their tastes in the population's distributions. The y-axis indicates the fraction of participants who answered the questions and the x-axis denotes the perceived rank of one's own position in the population's taste distribution. The original questions were asked over twenty points, but here they are grouped into five for clarity.

[cxlvi] A series of Anderson-Darling tests and Lilliefors tests were performed and suggests that, for all product categories, the distributions departed significantly from normality (W_{cheese} = 10.09, D_{cheese} = 0.22; $W_{\text{curry dishes}}$ = 9.78, $D_{\text{curry dishes}}$ = 0.22; W_{steak} = 7.71, D_{steak} = 0.17; $W_{\text{chocolate}}$ = 10.40, $D_{\text{chocolate}}$ = 0.21; W_{pillows} = 9.02, D_{pillows} = 0.19; $W_{\text{toothbrushes}}$ = 10.26, $D_{\text{toothbrushes}}$ = 0.22; $p < .0001$ for all tests). Likewise, the results of the Jarque-Bera test also reject the null of normally distributed population, while they produce slightly higher p-values for all product categories (JB_{cheese} = 16.36, p = 0.0040; $JB_{\text{curry dishes}}$ = 17.86, p = 0.0034; JB_{steak} = 11.48, p = 0.0105; $JB_{\text{chocolate}}$ = 10.15, p = 0.0140; JB_{pillows} = 11.73, p = 0.0097; $JB_{\text{toothbrushes}}$ = 14.15, p = 0.0060).

4.2.3 Accuracy of perceived relative standings in population's distribution and the role of knowledge

The normality (or non-normality) of the distributions of participants' subjective relative ranks of their tastes in the population's distributions is not sufficient to understand potential bias in people's estimates. To examine the accuracy of individual participants' estimations of their relative position in the population distribution, the prior ideal points were used as a proxy for the population's true tastes and the distribution of the ideal points provided by participants was treated as the relevant "true" distribution. Note that with the current dataset, there is no way to compute the exact rank positions for each participant since each ideal point is usually provided by more than one participant. Within each point, participants could not be ranked. Therefore, analysis could only identify minimum and maximum possible ranks for each participant[cxlvii]. The following analysis adopts two different, but complementary, approaches: a range approach and a mid-point approach. The former uses the computed minimum and maximum possible ranks as giving the range of the true rank, whereas the latter takes the middle point of this range.

The results obtained using the range approach indicate a moderate level of accuracy and an insignificant association between error size and self-reported knowledge of the product category. The percentage of participants whose subjective relative positions fell

---

[cxlvii] The minimum possible rank for a participant was computed by summing up the number of participants whose indicated values of prior ideal points were strictly lower than the target participant's. The maximum possible rank was calculated by subtracting the summation of the number of participants whose indicated values of prior ideal points were strictly higher than the target participant's from the number of participants who reported their prior ideal points. These estimated rank data were then transformed into percentage points by being divided by the number of participants who reported their prior ideal points. For example, suppose that there are 35 people who indicate their prior ideal points over a 7-point scale, and each point on the scale is chosen by the same number of people. Consider (as an example) people who indicate 3 as an ideal point. The minimum possible rank for them is $\frac{2\times5}{35}\times100\% = 28.57\%$, and the maximum possible rank is $\frac{35-(4\times5)}{35}\times100\% = 42.86\%$.

into the range of their "true" positions is 38.52% for cheese, 58.53% for curry dishes, 61.47% for steaks, 46.47% for chocolate, 66.40% for pillows, and 58.13% for toothbrushes. Overall, the error size for each product category is around 5%, with chocolate and cheese exhibiting the largest deviation, consistent with rates of accuracy. The average absolute deviation of participants' estimates from the true rank[cxlviii] is 0.0726 for cheese, 0.0519 for curry dishes, 0.0448 for steaks, 0.0960 for chocolate, 0.0475 for pillows, and 0.0478 for toothbrushes. Moreover, a Pearson's linear correlation analysis reveals that knowledge is not correlated with the error size at a 5% level of significance, regardless of whether the prior ideal points for each participant are controlled for. The only exception is chocolate. A negative correlation between participants' subjective level of knowledge and the error size was significant (r = -0.1278, p = 0.0475, CI = [-0.2501, -0.0014]). However, the correlation became insignificant when controlling for prior ideal points (r = -0.0107, p = 0.8690).

The mid-point approach found a higher level of accuracy and the rank order of accuracy among product categories was inconsistent with that in the previous case. Table 3 summarises the Pearson's linear correlations[cxlix] between the mid-point estimations of participants' true relative ranks and their subjective estimates of their relative positions in the population. As seen in the table, the correlations between participants' estimates and the computed true ranks were around 0.83 ± 0.05 and significant at the 5% level of significance. Among all product categories, cheese showed the highest estimated correlation coefficient, 0.8788. This was contrary to results obtained from the range approach, which indicated that the percentage of correct participants was lowest for the cheese category. Moreover, although participants appeared to perform best in the pillow

---

[cxlviii] This was computed by taking the difference between participants' perceived relative ranks and either maximum or minimum possible rank, depending on the direction of error.

[cxlix] The measure of accuracy here is different from the case of the range approach. Using a binary measure of accuracy for the mid-point approach may be misleading as it is almost impossible for the estimated relative rank to match the calculated true ranks, due to the elicitation methods used in the experiment. Indeed, according to the result, there is no participant whose mid-point true rank perfectly matches the self-estimated relative rank.

category in the previous case, the correlation for pillows was the second lowest among product categories when using the mid-point estimates. Yet, importantly, the accuracy observed using the mid-point approach was not significantly different between product categories, since most confidence intervals of correlations overlapped, as seen in the CIs presented in Table 3. Hence, while the rank orders of accuracy found using the two different approaches seemed to contradict each other, the results are not inconsistent[cl].

Table 3. Correlations of the "true" and the estimated relative rank positions in the population

Categories	*N*	*Corr*	*CI*
Cheese	244	0.8788	[0.8467, 0.9046]
Curry dishes	217	0.8436	[0.8004, 0.8781]
Steak	231	0.8521	[0.8124, 0.8840]

[cl] Nonetheless, the difference in calibration between two approaches reflects a methodological distinction between a point estimate and an interval. Owing to the limitation of computing exact estimated true ranks of participants, the range of ranks was wider for more frequently chosen ideal points than for unpopular ideal points. Put differently, if a participant whose indicated ideal point was most frequently chosen by other participants, the estimated range of true ranks for that participant would be wider than the range for any other participants who revealed other ideal points. Hence, the possibility of being inaccurate would be smaller than for others, regardless of whether participants themselves were biased in estimation. For example, in an extreme case in which all participants indicated the same prior ideal points, the results gained using the range approach *per se* would be meaningless. Yet, the mid-point estimate did not have this property. Contrarily, a wider range of ranks made the mid-point estimate less precise since the distance between any point within the range and the middle point would be farther when the range went larger. Therefore, participants who made a correct estimation (i.e., an estimation that was inside the range of possible true ranks) would be more likely to be considered biased and to exhibit a larger degree of inaccuracy if their indicated prior ideal points were more frequently chosen by others.

Chocolate	241	0.7856	[0.7320, 0.8296]
Pillows	247	0.8147	[0.7679, 0.8528]
Toothbrushes	246	0.8493	[0.8103, 0.8808]

Notes: All correlations presented are statistically significant at 1% level of significance, with all $p < .0001$. The upper and lower bounds of the confidence interval were computed at 5% level of significance, assuming data were normally distributed.

4.2.4 Compromise in choice

We now turn to the analysis of choices for each product category in a hypothetical shopping scenario. The aim is to examine whether a compromise effect is seen in the single-attribute methodology. As the survey did not measure choices in varying contexts, the existence of a tendency to compromise was measured through aggregate analysis of each choice distribution (focusing on mode, skewness, kurtosis, and absolute deviations from the middle point of the product line[cli], along with implementing three normality assessments like before, namely the Anderson-Darling test, the Jarque-Bera test, and the Lilliefors test). Note that this subsection also examines whether the choices were closer to or further from the middle of the consideration set than the ideal points were. This exploratory analysis is of particular importance since distributions of choice data and data on prior ideal points were alike in many product categories. In particular, both sets of data exhibited a marked tendency to concentrate around the middle point of the sets of options. Comparing these two patterns is helpful in understanding the origins of the observed compromise behaviour of choice data, if any. Furthermore, all analysis and tests were conducted without including participants who chose not to buy any product.

[cli] Absolute deviations from the middle were computed by transforming choice data in the way that the 4th product was coded as 0, the 3rd and 5th as 1, the 2nd and 6th as 2, and 1st and 7th as 3.

Figure 27 reveals a clear evidence of people's tendency to choose the 3rd or 5th product and to avoid the extreme options in the ordered set. This implies a possibility of a compromise tendency in a real choice decision. As seen in Figure 27, for categories such as cheese and chocolate, the 5th and 6th products were the most and the second most frequently chosen options respectively. The whole choice distributions for these two product categories were negatively skewed, with values of skewness and kurtosis respectively being -0.39 and 2.218 for cheese and -0.55 and 2.96 for chocolate. This suggests that although the shapes of cheese and chocolate's choice distributions were fairly symmetrical, since the values of skewness were within the range of ±1, the relatively large kurtosis values (> 2) show a high concentration of data points, indicating non-normality of the two choice distributions.

Conversely, for categories such as pillows and toothbrushes, products with ranks around the middle, i.e., the 3rd and 5th products, gained the highest choice share. Interestingly, the middle option, i.e., the 4th product, received many fewer choices. This result resembled the distributions of prior ideal points, such that the distributions of the middle three options were also in the shape of an inverse-U, with the 3rd or 5th options being the most popular over the whole set/scale. Moreover, Figure 27 clearly demonstrates that the choice distributions for pillows and toothbrushes were positively skewed, unlike those for cheese and chocolate. Specifically, the values of skewness and kurtosis respectively were 0.03 and 2.26 for pillows, but 0.15 and 1.99 for toothbrushes. Likewise, while the skewness of pillows' choice distribution was within the range to be considered normal, the distribution was slightly more heavy-tailed relative to a normal distribution. The category of toothbrushes, on the other hand, was found to have a normal choice distribution.

Choice data from the other two product categories, i.e., curry dishes and steak, were similar. Despite an opposite skewness, choice under these two categories was highly concentrated on just one product, either the 3rd or the 5th (see Figure 27). All other products were chosen far less often. Results for curry dishes resembled the results of prior ideal points, whereby almost 30% of participants indicated an ideal point of 5 out of 7, with other points acquiring only 10% of choices. Steak, however, did not show such similarity to prior ideal points. Surprisingly, although the choice distribution of steak could not be viewed as normal (skewness = 0.2988, K = 2.05), curry dishes' choice distribution (skewness = -0.25, K = 1.97) was considered to be normal since it was within the range of cut-off values.

Before moving towards formal tests to rigorously examine the normality of the distributions of choice data, the subsection first investigated how strongly choice data deviated from the middle points of the set in each product category, relative to data on prior ideal points. To this end, choice data and data on ideal points were transformed into absolute deviations (as stated in footnote cli). Overall, results were mixed for different product categories. First, deviations from the mid-point of the choice set were slightly greater in choice data, compared to data on prior ideal points, for categories such as curry dishes ($M_{choice} = 1.57$; $M_{ideal} = 1.56$), steak ($M_{choice} = 1.28$; $M_{ideal} = 1.26$), chocolate ($M_{choice} = 1.28$; $M_{ideal} = 1.22$), and toothbrushes ($M_{choice} = 1.40$; $M_{ideal} = 1.34$). Conversely, prior ideal points deviated more from the middle of the set than choice did for another two categories – cheese ($M_{choice} = 1.45$; $M_{ideal} = 1.47$) and ($M_{choice} = 1.20$; $M_{ideal} = 1.26$). However, a series of one-sample *t*-tests showed that none of the pairwise differences between data on deviation are significant ($t_{cheese}(243) = -0.6968$, $p_{cheese} = .4866$; $t_{curry\ dishes}(207) = 0.2127$, $p_{curry\ dishes} = .8318$; $t_{steak}(228) = 0.6615$, $p_{steak} = .5090$; $t_{chocolate}(238) = 1.3806$, $p_{chocolate} = .1687$; $t_{pillows}(246) = -1.5599$, $p_{pillows} = .1201$; $t_{toothbrushes}(244) = 1.4793$, $p_{toothbrushes} = .1403$). This implies a potential influence of (biased) prior ideal points[clii] on choice of products, which was probably due to the influence of self-reported preferences on downstream choice and/or the existence of common underlying factors, such as social norms, that affect both data and/ or similarity in their elicitation methods[cliii].

[clii] For more details, please refer to footnote cxlii.

[cliii] Both choice and prior ideal points were measured using a seven-point scale. Although the experiment tried to minimise the effect of answers of prior ideal points on choice by separating these two sets of questions from each other in the experiment, it is possible that participants still remembered their answers on prior ideal point questions and applied them directly to choice questions.

Finally, as before, a more formal normality assessment on choice distribution was conducted using the Anderson-Darling test, the Jarque-Bera test, and the Lilliefors test[cliv]. Results suggested that choice distributions significantly differed from normality for all categories, except pillows. To conclude, with the degree of deviation from the middle point of the choice sets similar to that of prior ideal points', data collected from the experiment overall supported the notion of an inclination to compromise in choice for a wide range of product categories, but most of the choice distributions were not perfectly bell-shaped or statistically normal.

[cliv] Results obtained from a series of Anderson-Darling tests and Lilliefors tests indicate a rejection of the null hypothesis that a choice data came from a normal distribution (W_{cheese} = 7.64, D_{cheese} = 0.21; $W_{\text{curry dishes}}$ = 6.57, $D_{\text{curry dishes}}$ = 0.22; W_{steak} = 7.72, D_{steak} = 0.22; $W_{\text{chocolate}}$ = 8.08, $D_{\text{chocolate}}$ = 0.228; W_{pillows} = 6.48, D_{pillows} = 0.18; $W_{\text{toothbrushes}}$ = 6.73, $D_{\text{toothbrushes}}$ = 0.17; $p < .0001$ for all tests). Likewise, the Jarque-Bera test for all product categories, except pillows, also rejected the null that each distribution stemmed from a normal distribution with an unknown mean and variance, while this test produced higher p-values (JB_{cheese} = 12.47, p = 0.0083; $JB_{\text{curry dishes}}$ = 11.41, p = 0.0109; JB_{steak} = 11.95, p = 0.0095; $JB_{\text{chocolate}}$ = 12.03, p = 0.0092; $JB_{\text{toothbrushes}}$ = 11.37, p = 0.0106). Regarding pillows, there is no sufficient evidence to reject the null at 5% level of significance (JB_{pillows} = 5.68, p = 0.0510), implying a normality. This result is in line with Yap and Sim's (2011) power comparison, which suggests that the Jarque-Bera test works better for symmetric long-tailed distributions than other tests do, whereas the Anderson-Darling test is more powerful for asymmetric distributions. The shape of pillows' choice distribution displayed in Figure 27 shows that the Jarque-Bera test is more suitable and detectable than the Anderson-Darling test is.

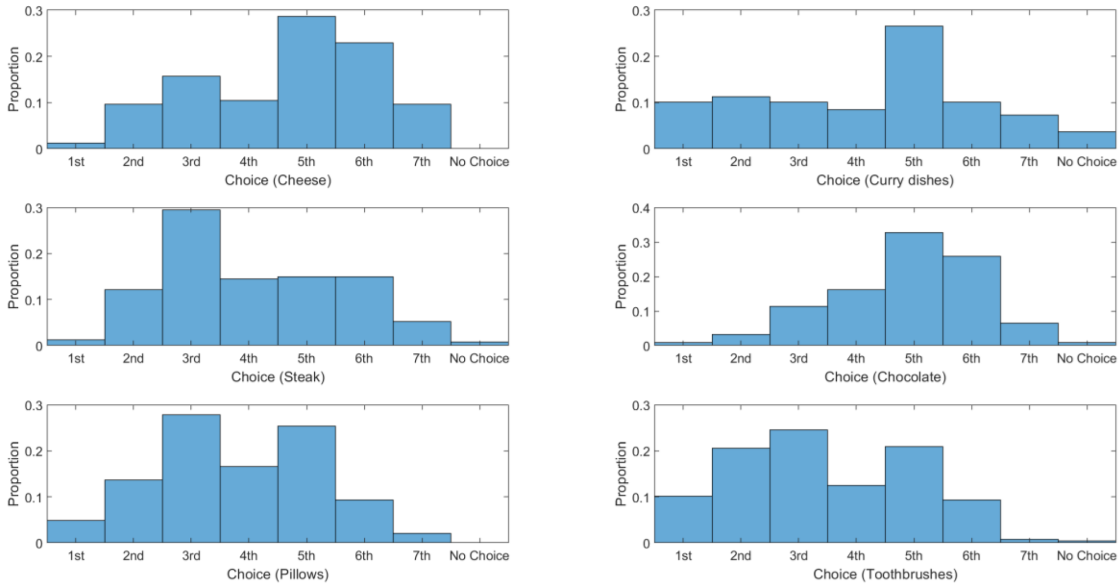

Figure 27. Relative choice frequency of products for each category. The y-axis refers to the fraction of participants who chose the product and the x-axis denotes relative ranks of products. The "No Choice" bin represents the proportion of participants who chose to not to purchase a product. Its relatively high fraction in the category of curry dishes reflects the fact that many participants were uncertain about their preference towards listed products. This finding echoes an observation of a low self-expressed level of knowledge about curry spiciness as well as a high percentage of participants who did not indicate their prior ideal points for curry spiciness, although those participants had already been excluded from seeing this question.

4.2.5 Effects of knowledge on the compromise effect

It has been suggested in the literature that a lack of knowledge about a product category may be a driver of the compromise effect (Sheng, Parker, & Nakamoto, 2005). This subsection therefore investigates the association between knowledge and a tendency to choose the middle option. To this end, choice data was, again, transformed into absolute deviation from the middle. Data on choosing nothing was not included in the correlation analysis. The analysis was conducted with a Pearson's linear correlation. If the hypothesis about the role of knowledge in reducing the compromise effect is true, the estimated correlations between deviation (from middle) and knowledge would be positive.

Overall, results were mixed for different product categories, with more running counter to the hypothesis. A Pearson's linear correlation showed that for product categories like cheese, steak, pillows, and toothbrushes, knowledge is not significantly correlated with the tendency to avoid compromise options (r_{cheese} = 0.1204, p_{cheese} = 0.0604; r_{steak} = 0.0902, p_{steak} = 0.1739; $r_{pillows}$ = 0.0902, $p_{pillows}$ = 0.1351; $r_{toothbrushes}$ = 0.1239, $p_{toothbrushes}$ = 0.0527), though most of them were significant at the 10% level of significance and all had positive signs. Results for product categories such as curry dishes and chocolate, in contrast, reveal a significant negative linear association ($r_{curry\ dishes}$ = -0.1467, $p_{curry\ dishes}$ = 0.0345; $r_{chocolate}$ = 0.2069, $p_{chocolate}$ = 0.0013). This correlation, however, becomes insignificant[clv] after controlling for prior ideal points, which were also transformed into absolute distance from the middle[clvi] ($r_{curry\ dishes}$ = -0.1224, $p_{curry\ dishes}$ = 0.0789; $r_{chocolate}$ = 0.0575, $p_{chocolate}$ = 0.3772). The finding of an insignificant partial linear correlation suggests that knowledge might affect choice via its association with prior ideal points. This speculation was partially supported by observed significant (partial) correlations between absolute deviation from the middle in data of prior ideal points and that in choice[clvii], regardless of whether knowledge is included as a control, as well as a significant association between knowledge and absolute departure from the middle in prior ideal points for chocolate ($r_{chocolate}$ = 0.2434, $p_{chocolate}$ = .0001), but not for curry dishes ($r_{curry\ dishes}$ = -0.1224, $p_{curry\ dishes}$ = 0.0789).

[clv] Insignificant linear associations were also observed for other categories (r_{cheese} = 0.0064; p_{cheese} = 0.9213; r_{steak} = -0.0037; p_{steak} = 0.9555; $r_{pillows}$ = 0.0148; $p_{pillows}$ = 0.8169; $r_{toothbrushes}$ = 0.1153; $p_{toothbrushes}$ = 0.0723).

[clvi] The transformation method is same as the one for choice data.

[clvii] The correlation analysis was conducted with the transformed data as before. The correlation coefficient was 0.7770 for cheese, 0.7586 for curry dishes, 0.8163 for steak, 0.7252 for chocolate, 0.7027 for pillows, and 0.7390 for toothbrushes, with all having $p < .0001$. When controlling for knowledge, the partial correlation was still significant ($p < .0001$). The partial correlation coefficient for the same order of product categories respectively was 0.7732, 0.7568, 0.8146, 0.7116, 0.6995, and 0.7383.

Again, integrating findings on Pearson's linear correlation under the category of chocolate, it seems that knowledge might decrease the compromise tendency in choice through its positive link to absolute deviation (from the middle) in prior ideal points. In other words, the results imply that people who believe themselves to have a higher level of knowledge (versus those who state a lower level) may be more confident about their preferences and thus show less compromise when indicating prior ideal points[clviii], which is reflected in a positive correlation between knowledge and a tendency to reveal more extreme ideal points.

The results for curry dishes, however, reveal a completely different story, reducing generalisability of the above conjecture. Firstly, the negative correlation between knowledge for curry dishes and tendency to choose more extreme options suggests that decision makers who believed that they are highly knowledgeable about curry dishes would be more likely to choose products placed at the middle of the choice set, contradicting the hypothesis that knowledge attenuates the compromise effect. In addition, the insignificant (negative) relationship between knowledge and a tendency to indicate more extreme ideal points goes against the conjecture that prior ideal points act as a meditator variable. This result may be related to the observation that the partial correlation between knowledge and a choice of products ranked close to an extreme of the consideration set, after controlling for ideal points, still approaches significance (reaching the 10% level). All of this evidence indicates that there may be a direct, positive association between knowledge and a compromise tendency in choice of curry dishes, inconsistent with findings in the chocolate case.

Overall, the evidence is equivocal, with only little evidence being found to support the hypothesis about the mitigating role of knowledge in the compromise effect.

[clviii] It is important to note that data on prior ideal points gathered from the experiment is, by definition, not preferences$_U$, which is not supposed to be influenced by one's knowledge.

4.2.6 Beliefs about a rational design of a product line

Before turning to the evidence of the rank-order decision rule, it is helpful to investigate a crucial premise first. This is the idea that people typically believe that product lines are designed to capture the whole population's tastes and also think there is no menu distortion. If the first part of the premise does not hold, there will be no reason for the rank order decision rule to be used during decision making. However, if it is the second part that is invalid, the rank order decision rule may be adopted, but not in a direct way (i.e., individuals' estimates of market demand for products will be involved in the matching process). In other words, when menu is thought to be distorted, the possibility for a decision maker to employ a "direct" matching strategy will be small as it is less useful. Recall that the direct matching strategy refers to a process whereby a decision maker directly matches his/ her estimate of their relative position in the population with respect to tastes to a product's relative rank in an ordered set[clix]. This contrasts with a more comprehensive matching strategy which also takes into account a non-uniform distribution of population choice. A direct comparison in choice predictability between the direct matching approach and the more comprehensive version will be presented in the next subsection.

Data collected from questions that asked participants to estimate other consumers' choices by assigning twenty tokens to available products were analysed and two necessary conditions (that need to be met if the direct matching strategy is to be assessed) were considered. First, if participants believed the design of market offerings accurately reflected the population's taste distribution, they would assign at least one token to every product provided. Second, distribution of market choice estimated by a participant should be (at least asymptotically) uniform. A non-flat estimated distribution of market choice may bias precision of choice predicted using direct matching. For example, if a decision maker perceives himself/ herself to be located at the 50th percentile of the population's

---

[clix] For instance, if a person believes he or she is at 40% of the population's distribution in terms of preferences for cheese sharpness, he or she will buy a cheese with a relative rank of 40%, computed based on its degree of sharpness relative to other cheese in the market.

taste distribution, and believe products ranked the 1st and the 2nd in the set equally generate 30% of market demand, with the remaining five options in the set sharing 40% of demand, the decision maker will choose the 2nd product, according to the rank order decision rule. This choice outcome, however, is inconsistent with a direct matching strategy, which suggests that the decision maker should choose the middle (4th) option in the set. Predictions of individual choice based on a direct matching method will thus be misleading.

The results only supported the first condition. The percentage of participants who assumed that all products would be chosen by at least 5% of the population was 85.66% for cheese, 82.49% for curry dishes, 76.62% for steak, 82.16% for chocolate, 82.19% for pillows, and 80.08% for toothbrushes. This a high proportion of participants believed that no product would attract zero demand in the market. However, the second condition for using a direct matching strategy was not satisfied. Data showed that no participant assigned fairly equal amounts of tokens (i.e., 2 to 3 tokens) to all products for all product categories[clx]. In fact, the middle option, on average, received more tokens than any other options ($M_{\text{cheese, 4th}} = 3.86$, $SD_{\text{cheese, 4th}} = 2.07$; $M_{\text{curry dishes, 4th}} = 3.54$, $SD_{\text{curry dishes, 4th}} = 1.92$; $M_{\text{steak, 4th}} = 4.57$, $SD_{\text{steak, 4th}} = 2.58$; $M_{\text{chocolate, 4th}} = 4.20$, $SD_{\text{chocolate, 4th}} = 2.60$; $M_{\text{pillows, 4th}} = 4.17$, $SD_{\text{pillows, 4th}} = 2.43$; $M_{\text{toothbrushes, 4th}} = 4.48$, $SD_{\text{toothbrushes, 4th}} = 2.90$). The extreme options, in contrast, gained fewest tokens ($M_{\text{cheese, 1st}} = 2.50$, $SD_{\text{cheese, 1st}} = 1.61$; $M_{\text{curry dishes, 1st}} = 3.05$, $SD_{\text{curry dishes, 1st}} = 2.78$; $M_{\text{steak, 1st}} = 1.44$, $SD_{\text{steak, 1st}} = 1.36$; $M_{\text{chocolate, 1st}} = 1.98$, $SD_{\text{chocolate, 1st}} = 1.71$; $M_{\text{pillows, 1st}} = 2.72$, $SD_{\text{pillows, 1st}} = 2.41$; $M_{\text{toothbrushes, 1st}} = 2.85$, $SD_{\text{toothbrushes, 1st}} = 2.31$; and $M_{\text{cheese, 7th}} = 2.22$, $SD_{\text{cheese, 7th}} = 1.77$; $M_{\text{curry dishes, 7th}} = 2.12$, $SD_{\text{curry dishes, 7th}} = 2.01$; $M_{\text{steak, 7th}} = 2.07$, $SD_{\text{steak, 7th}} = 2.25$; $M_{\text{chocolate, 7th}} = 2.15$, $SD_{\text{chocolate, 7th}} = 2.24$; $M_{\text{pillows, 7th}} = 1.95$, $SD_{\text{pillows, 7th}} = 1.89$; $M_{\text{toothbrushes, 7th}} = 1.48$, $SD_{\text{toothbrushes, 7th}} = 1.24$). Many of the pairwise differences are statistically significant, according to *t* tests. This suggests that participants in general believed that market demand for products is not equal, with the middle option

[clx] For each product category, there were 20 tokens but 7 products to be assigned to. Therefore, uniformity of the token allocation requires participants to assign 2 or 3 tokens to each product.

being more popular than others. Accordingly, survey results reveal that although the use of a rank order decision rule is highly likely, decision makers may not do it via a direct matching.

To conclude, results partially supported the requirements for a rank order matching rule to be used. Specifically, most of participants thought product lines were designed to contain a population's taste distribution, although estimated market demand over products was not uniformly distributed. Therefore, it is possible that participants did not directly match their estimated relative standing in taste distribution to a product's relative rank to make a purchasing decision. This in turn implies that the false consensus effect per se is not sufficient to explain the compromise effect, even when a rank order decision rule is used. It is concluded that any choice prediction made from the rank order decision rule needs to embody people's own estimation about a market's choice distribution over products.

4.2.7 A direct versus comprehensive rank order decision rule

As found in the previous subsection, a direct matching process may not occur during decision making since participants normally did not expect a flat market demand over market offerings. Thereby, a matching strategy must incorporate decision makers' own estimates of choice distributions in the market. To further investigate this question, this subsection explores choice predicted by two different implementations of a rank order decision rule, and examines the accuracy of their predictions. The first method is the direct approach, which directly matches people's perceived relative standing in the population's distribution of a certain taste on an attribute to a product's relative position in the set, ranked in terms of the given attribute[clxi]. In other words, individuals who adopt a direct matching strategy will choose a product whose relative rank in the menu is identical to their self-estimated relative standings in the population. The second method is referred to

---

[clxi] Computationally, this was achieved by rescaling a twenty-point scale (individuals' estimation about their relative rank in the population) into a seven-point scale (a product choice set).

as a comprehensive approach, which additionally accommodates decision makers' own estimation about a distribution of a market demand into the rank order decision rule. For example individuals who believe that their preferences lie at the median of the population distribution may choose the second market option from a set of seven products if they also believe that population preferences and hence choices are distributed over market offerings in a positively skewed way.

Overall, results revealed a difference between choice predicted by the two approaches. As shown in Figure 28, choice predicted by the comprehensive approach was less skewed than what is predicted by the direct approach, because market demand on available products, on average, was thought to be bell-shaped. This implies that the false consensus effect in self-estimated relative positions in the population's preference distribution may not be the sole contributor to the compromise effect. The rather symmetrical distributions of market demand estimated by individual consumers might strengthen their tendency to choose the compromise product, if the comprehensive version of the rank order decision rule was adopted. Moreover, a series of *t*-tests also found differences in predicted choice between the two approaches for many product categories. Predicted choice differed significantly between estimation approaches for categories, including cheese ($M_{comprehensive} = 4.51$ vs. $M_{direct} = 4.62$; $t(243) = -2.36$, $p = .0189$), curry dishes ($M_{comprehensive} = 3.94$ vs. $M_{direct} = 4.06$; $t(216) = -2.26$, $p = .0246$), pillows ($M_{comprehensive} = 3.61$ vs. $M_{direct} = 3.83$; $t(246) = -3.99$, $p < .0001$), and toothbrushes ($M_{comprehensive} = 3.28$ vs. $M_{direct} = 3.50$; $t(245) = -4.76$, $p < .0001$). Yet, the difference was insignificant for categories like steak ($M_{comprehensive} = 4.18$ vs. $M_{direct} = 4.12$; $t(230) = 1.15$, $p = .2495$) and chocolate ($M_{comprehensive} = 4.76$ vs. $M_{direct} = 4.78$; $t(240) = -0.51$, $p = .6105$).

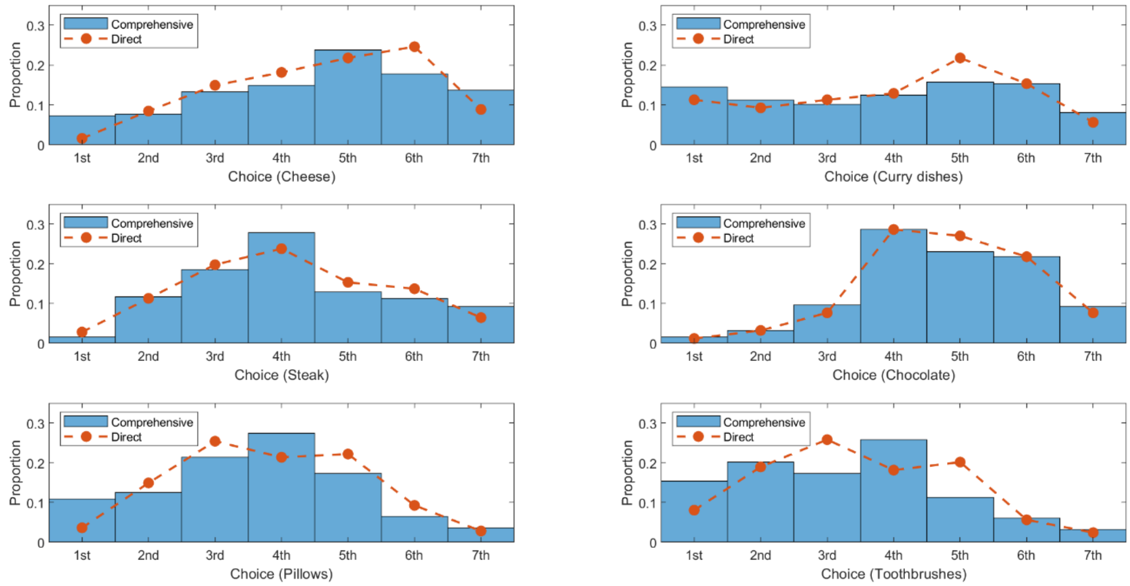

Figure 28. Choice share predicted using the direct matching and the comprehensive approach. As the rank order decision rule was not able to predict the probability of opting out, choice share for the "No Choice" option was not computed. Examination of the figures reveals that choice share predicted by the direct approach essentially echoes Figure 26 because it was directly transformed from the data of self-estimation about one's relative standing in the population's taste distribution. In contrast, choice share predicted using the comprehensive approach was more symmetrically distributed, with a centre located at the middle. This is because the comprehensive approach considered participants' estimation about distributions of market demand, which was highly bell-shaped.

Furthermore, relative accuracy of choice estimates using the two approaches was explored[clxii]. Surprisingly, the direct matching method performed much better than the comprehensive approach did, in terms of their predictability of actual choice. A *t*-test showed that the difference in the average rates of accuracy[clxiii] between the two

---

[clxii] Note that since the rank order decision rule could not predict the situation of opting out, the following analysis excluded data from participants who chose nothing from the set.

[clxiii] The average rates of accuracy refers to the rate of accuracy averaged across participants. In practice, the number of participants whose actual choice was identical to

approaches was significant ($M_{comprehensive} = 0.53$ vs. $M_{direct} = 0.65$, $t(5) = -6.93$, $p < .0001$). Collapsing across product category, the predictions from the direct method were still more accurate than those from the comprehensive approach for most of the categories ($M_{comprehensive, cheese} = 0.50$ vs. $M_{direct, cheese} = 0.62$, $t_{cheese}(243) = -3.17$, $p_{cheese} = 0.0017$; $M_{comprehensive, curry dishes} = 0.51$ vs. $M_{direct, curry dishes} = 0.64$, $t_{curry dishes}(216) = -3.28$, $p_{curry dishes} = 0.0012$; $M_{comprehensive, chocolate} = 0.55$ vs. $M_{direct, chocolate} = 0.65$, $t_{chocolate}(240) = -2.65$, $p_{chocolate} = 0.0085$; $M_{comprehensive, pillow} = 0.47$ vs. $M_{direct, pillow} = 0.67$, $t_{pillow}(246) = -4.91$, $p_{pillow} < .0001$; $M_{comprehensive, toothbrushes} = 0.57$ vs. $M_{direct, toothbrushes} = 0.71$, $t_{toothbrushes}(245) = -4.10$, $p_{toothbrushes} < .0001$). The only exception is steak, where there was no significant difference in mean rate of accuracy ($M_{comprehensive, steak} = 0.57$ vs. $M_{direct, steak} = 0.62$, $t_{steak}(230) = -1.59$, $p_{steak} = 0.1124$).

Similarly, the correlation between participants' actual choice and predicted choice suggests that estimates obtained with the direct method have a higher correlation coefficient for every product category ($r_{cheese} = 0.8978$; $r_{curry dishes} = 0.8671$; $r_{steak} = 0.8676$; $r_{chocolate} = 0.8148$; $r_{pillows} = 0.8489$; $r_{toothbrushes} = 0.9095$; all $p < .0001$), relative to the comprehensive approach ($r_{cheese} = 0.8476$; $r_{curry dishes} = 0.8171$; $r_{steak} = 0.8400$; $r_{chocolate} = 0.8000$; $r_{pillows} = 0.7814$; $r_{toothbrushes} = 0.8903$; all $p < .0001$). However, the differences in estimated correlations are not significant, given that their 95% confidence intervals[clxiv] overlap each other for each product category.

---

predicted choice for each product was first computed. Then, each number was divided by the number of participants who chose from the menu provided. This gave a vector of accuracy rates for each product category, i.e., a 1-by-6 vector. The t-tests were conducted on two vectors computed using the direct and comprehensive approach respectively.

[clxiv] The lower and higher bounds of the confidence intervals of the estimated correlations were as follows. In the case of the comprehensive approach, *CI* = [0.8080, 0.8796] for cheese, [0.7663, 0.8578] for curry dishes, [0.7971, 0.8744] for steak, [0.7492, 0.8415] for chocolate, [0.7276, 0.8257], for pillows, and [0.8610, 0.9137] for toothbrushes, whereas in the case of the direct approach, *CI* = [0.8703, 0.9197], [0.8288, 0.8973], [0.8314, 0.8964], [0.7672, 0.8534], [0.8100, 0.8805], and [0.8850, 0.9289], respectively.

Moreover, Table 4 displays conditional frequencies of predicted choice, given actual choice, for each product category. In other words, each cell in the table gives the percentage of participants who were expected to choose a *n*-th product (column), based on the fact that a particular product (row) was actually chosen in the experiment. Consistent with previous results, Table 4 revealed that the direct method is, in general, closer to participants' actual performance than the comprehensive method. This was further confirmed by a *t*-test on conditional frequencies of correct predictions[clxv] (i.e., the number of predicted choices being the same as a given actual choice, relative to the number of given actual choices being mismatched to predicted choice), averaged across all seven products, under each product category ($M_{comprehensive}$ = 59.16% vs. M_{direct} = 67.00%, $t(5) = -3.73$, $p = 0.0136$). More interestingly, Table 4 shows that for less extreme-ranked products, e.g., the 3rd, 4th, and 5th product, the direct method were far more accurate than the comprehensive approach, compared to results for extreme-ranked products. One plausible, ad-hoc explanation is that choices predicted by the comprehensive approach were more concentrated on the middle of the choice set (see Figure 28). This may lead to a distortion of the predicted possibility of choosing the middle products.

In summary, experimental evidence suggests that consumers may use a rank order decision rule when making a purchase decision, but that the way and/or reason they do it seems not to follow what Wernerfelt (1995) asserted. More specifically, according to Wernerfelt (1995), the rationale of using the rank order decision rule is that consumers know more about their tastes compared to others, and normally believe that menu unbiasedly convey population's taste-relevant information. However, the present experiment findings indicate that this taste-relevant information appears to be trivial when consumers adopt this rule. In other words, people seem to directly match their relative standings in the distribution in terms of tastes to ranks of products in the market, without considering their beliefs about the population's choice distribution over products. Yet, it

---

[clxv] The conditional frequencies of correct prediction refer to numbers in a bold form in Table 4.

is also important to note that, as suggested in Model 1, the rank order decision rule may not drive choice alone – the final choice may be a result of different forces acting against each other. Thus, the observed inconsistency with the model assumptions may stem from the fact that there are other drivers of choice.

Table 4. Conditional proportion of predicted choice given actual choice for each product category.

	The comprehensive approach (%)							The direct approach (%)						
Real Choice	1st	2nd	3rd	4th	5th	6th	7th	1st	2nd	3rd	4th	5th	6th	7th
Cheese														
1st	**100.00**	0	0	0	0	0	0	**66.67**	33.33	0	0	0	0	0
2nd	54.17	**37.50**	4.17	0	4.17	0	0	8.33	**62.50**	20.83	8.33	0	0	0
3rd	5.13	17.95	**53.85**	12.82	7.69	2.56	0	0	12.82	**61.54**	23.08	2.56	0	0
4th	0	7.69	23.08	**50.00**	19.23	0	0	0	0	23.08	**65.38**	11.54	0	0
5th	0	0	7.04	19.72	**47.89**	22.54	2.82	0	0	2.82	21.13	**56.34**	18.31	1.41
6th	0	1.75	0	7.02	26.32	**42.11**	22.81	0	0	0	3.51	17.54	**68.42**	10.53
7th	0	0	0	4.17	4.17	12.50	**79.14**	0	0	0	0	0	37.5	**62.50**
Curry dishes														
1st	**80.00**	12.00	8.00	0	0	0	0	**84.00**	12.00	0	4.00	0	0	0
2nd	17.86	**46.43**	14.29	14.29	0	3.57	3.57	7.14	**50.00**	21.43	10.71	7.14	3.57	0
3rd	16.00	28.00	**36.00**	12.00	4.00	4.00	0	8.00	20.00	**52.00**	16.00	0	4.00	0
4th	4.76	4.76	33.33	**47.62**	0	9.52	0	0	4.76	19.05	**57.14**	14.29	4.76	0
5th	0	0	4.55	18.18	**48.48**	27.27	1.52	0	0	1.52	7.58	**72.73**	18.18	0

	The comprehensive approach (%)							The direct approach (%)						
Real Choice	1st	2nd	3rd	4th	5th	6th	7th	1st	2nd	3rd	4th	5th	6th	7th
6th	0	8.00	0	8.00	16.00	**52.00**	16.00	0	0	4.00	12.00	4.00	**72.00**	8.00
7th	0	5.56	0	0	5.56	11.11	**77.78**	0	0	5.56	5.56	0	22.22	**66.67**
Steak														
1st	**100.00**	0	0	0	0	0	0	**100.00**	0	0	0	0	0	0
2nd	3.33	**70.00**	20	6.67	0	0	0	13.33	**70.00**	13.33	3.33	0	0	0
3rd	0	9.59	**47.95**	36.99	4.11	0	1.37	0	9.59	**54.79**	28.77	5.48	1.37	0
4th	0	0	5.56	**72.22**	19.44	2.78	0	0	0	2.78	**83.33**	8.33	5.56	0
5th	0	2.70	2.70	24.32	**40.54**	21.62	8.11	0	0	2.70	13.51	**59.46**	24.32	0
6th	0	0	2.70	13.51	18.92	**48.65**	16.22	0	0	5.41	5.4	18.92	**51.35**	18.92
7th	0	0	0	0	0	0	**100.00**	0	0	0	0	7.69	23.08	**69.23**
Chocolate														
1st	**50.00**	0	50.00	0	0	0	0	**50.00**	0	0	50.00	0	0	0
2nd	37.50	**50.00**	12.50	0	0	0	0	25	**62.50**	0	12.50	0	0	0
3rd	0	14.29	**53.57**	25.00	3.57	3.57	0	0	10.71	**60.71**	25.00	3.57	0	0
4th	0	0	7.500	**80.00**	10.00	0	2.50	0	0	2.50	**90.00**	2.50	5	0
5th	0	0	2.47	33.33	**43.21**	19.75	1.23	0	0	1.23	18.52	**64.20**	14.81	1.23
6th	0	0	1.56	6.25	26.56	**51.56**	14.06	0	0	0	12.50	18.75	**56.25**	12.5
7th	0	0	0	6.25	0	18.75	**75.00**	0	0	0	6.25	6.25	25	**62.50**
Pillows														
1st	**66.67**	25	8.33	0	0	0	0	**58.33**	33.33	0	0	8.33	0	0

Real Choice	The comprehensive approach (%) 1st	2nd	3rd	4th	5th	6th	7th	The direct approach (%) 1st	2nd	3rd	4th	5th	6th	7th
2nd	41.18	**38.24**	17.65	0	2.94	0	0	5.88	**64.71**	23.53	2.94	2.94	0	0
3rd	5.80	21.74	**50.72**	20.29	1.45	0	0	0	14.49	**68.12**	17.39	0	0	0
4th	0	0	12.20	**68.29**	12.20	2.44	4.88	0	0	12.20	**75.61**	12.20	0	0
5th	0	0	6.35	36.51	**42.86**	14.29	0	0	0	1.59	12.70	**68.25**	17.46	0
6th	4.35	0	4.35	8.70	39.13	**26.09**	17.39	0	4.35	4.35	4.35	21.74	**52.17**	13.04
7th	0	0	20.00	20.00	0	0	**60.00**	0	0	20.00	0	0	0	**80.00**
Toothbrushes														
1st	**76.00**	20.00	4.00	0	0	0	0	**64.00**	24.00	8.00	4.00	0	0	0
2nd	35.29	**54.90**	7.84	1.96	0	0	0	7.84	**66.67**	21.57	3.92	0	0	0
3rd	1.64	24.59	**54.10**	19.67	0	0	0	0	11.48	**73.77**	13.11	1.64	0	0
4th	0	3.23	16.13	**77.42**	3.23	0	0	0	0	16.13	**77.42**	6.45	0	0
5th	0	0	0	46.15	**44.23**	9.62	0	0	0	1.92	13.46	**80.77**	3.85	0
6th	0	4.35	0	8.70	17.39	**43.48**	26.09	0	0	0	8.70	21.74	**52.17**	17.39
7th	0	0	0	0	0	0	**100.00**	0	0	0	0	0	0	**100.00**

Notes: Each raw under each approach should sum up to 100 (%). Numbers highlighted in bold refers to a percentage of predicted choice being same as actual choice. It therefore implies rates of accuracy of choice predictions.

4.2.8 Preferences$_E$

This subsection explores participants' preferences$_E$ to determine whether stated preferences were consistent with choice[clxvi]. As described in the Method section, data about preferences$_E$ were collected by asking participants to rate their preference for each available product on a 5-point scale, from "Not at all" to "Like very much". These data were transformed into a numeric array, whereby "Not at all" was coded as 1 and "Like very much" was coded as 5.

For all product categories, results show that preferences$_E$ strongly reflect choice. First, it was found that products that were chosen in the hypothetical purchase stage received, on average, higher ratings than other products ($M_{cheese, choice} = 4.55$, $M_{cheese, others} = 3.48$; $M_{curry dishes, choice} = 4.40$, $M_{curry dishes, others} = 3.03$; $M_{steak, choice} = 4.56$, $M_{steak, others} = 2.65$; $M_{chcoclate, choice} = 4.51$, $M_{chocolate, others} = 3.29$; $M_{pillows, choice} = 4.44$, $M_{pillows, others} = 2.86$; $M_{pillows, choice} = 4.43$, $M_{pillows, others} = 2.67$). Similarly, of participants who chose to buy a product in the hypothetical supermarket, many rated their chosen products higher than other products, with the exact proportion of participants being 89.79% for cheese, 89.90% for curry dishes, 93.45% for steak, 92.89% for chocolate, 87.04% for pillows, and 92.65 for toothbrushes. Note that these numbers include participants who rated chosen products equally highly to other products. Indeed, the fraction of participants who rated chosen products as the only favourite one was much lower (39.34% for cheese, 45.19% for curry dishes, 58.08% for steak, 37.24% for chocolate, 51.42% for pillows, and 57.96% for toothbrushes). This implies that participants might express more than one favourite products, but could nevertheless choose one if necessary.

4.3 Summary

This chapter aimed to explore the main assumptions of the models proposed in Chapters 2 and 3. An online experiment was conducted to investigate the models' assumptions and predictions. In general, the experiment results support the hypotheses

---

[clxvi] Preferences$_E$, by definition, are the preferences people reveal to others, often through choice. This preference is, therefore, expected to reflect choice.

concerning the assumptions and implications of the models. However, it is important to note that due to technical limitations of any online experiment, such as the inability to measure preferences$_U$ or to offer participants real products to experience, the overall performance of the proposed models cannot be assessed directly.

As predicted, the assumption of the false consensus effect in preferences is supported by the experimental evidence. The results clearly indicate a tendency for participants to believe that they stand close to the middle of the population's taste distribution, reflecting the fact that people normally believe that they are rather average compared to other consumers in terms of their preferences for a product attribute level. However the analysis only focuses on the implications of the false consensus effect, instead of directly proving the effect itself, since the proposed models do not incorporate the false consensus effect with its original meaning. The measurement of the false consensus effect in the present experiment is different from the traditional paradigm, in that classic experiments typically measure it by assessing a relative difference between estimates made by people who support an opinion or behaviour and estimates made by opponents (Marks & Miller, 1987; Mullen & Hu, 1988; Mullen, Dovidio, Johnson, & Copper, 1992). That is, conventionally, the false consensus effect is not estimated by asking participants about their relative positions in the population.

Despite the tendency for people to estimating themselves to be less extreme than other people in terms of tastes about a product attribute, people's estimates about their relative positions in the population were only moderately accurate for all product categories. This evidence is consistent with prior research, which suggest that people's estimates of their relative standings are often error-prone and susceptible to various environmental cues (Burson, 2007; Burson & Gershoff, 2015; Gershoff & Burson, 2011). This estimation error is typically attributed to people's inability to correctly assess absolute magnitudes of their characteristics such as ability and preferences (Freund & Kasten, 2012; Krueger & Mueller, 2002; Zell & Krizan, 2014) or a cognitive tendency to overestimate the dispersion of other people (Gershoff & Burson, 2011). Importantly, the majority of previous research on this topic is conducted in the domain of ability and knowledge, which is advantageous in its objectively and subjectively assessable nature. The present study, however, revolves around preferences, where objective values do not exist in the same sense. Recall that the analysis uses prior ideal points, i.e., participants' self-reported preferences, to compute participants' "true" relative positions in the population's taste distribution. This measure is admittedly less precise than objectively assessing knowledge or ability via a well-designed test, since people's self-assessed levels of preferences may, like their estimation about own relative standings, be biased[clxvii]. Lack of an objective measure of true levels of preferences may be one reason why observed accuracy of perception about one's own relative positions in the population's tastes distribution seems better than previous findings on knowledge and ability as well as why no significant association between levels of knowledge and sizes of estimation error is observed.

Moreover, results obtained from the present experiment show a marked tendency for participants to select middle products and avoid extreme alternatives in an ordered menu. This implies that a compromise tendency in purchase decisions can be found in a case where products in a given categories are differentiated in terms of a single attribute and a no-choice opportunity is offered in choice tasks. In addition, the extent to which choice data deviated from its middle point (i.e., 4) is not statistically different from data on the tendency of prior ideal points to do so. This suggests that self-reported prior ideal points may have some influences on final choice, through either a pathway of preferences shaping choice or common, external underlying factors such as social norms, or both. This finding may also be due simply to similarity in elicitation methods of both data.

Note that, as with the investigation of the false consensus effect, the analytical methods employed to examine the compromise effect differ considerably from conventional paradigms. In particular, classical approaches typically measure the

---

[clxvii] This argument is in line with the present finding that distributions of prior ideal points are approximately normal and similar to distributions of choice data. Of course, this evidence is not enough to conclude that there must be a systematic bias in estimating prior ideal points, given that many human traits in population are normally distributed. Yet, the possibility that there exist external forces, such as social norms, that affect evaluations of ideal points also cannot be ruled out since the nature of prior ideal points is more like preferences$_I$ or expressed preferences, which are possibly subject to external influences, in addition to preferences$_U$.

compromise effect by computing changes in choice share of a non-middle option after it becomes a middle option as a result of the addition of a more extreme alternative to the pre-existing set (e.g., Sheng, Parker & Nakamoto, 2005; Simonson & Tversky, 1992) and study it with multi-dimensional products (e.g., Drolet, Luce, & Simonson, 2008; Khan, Zhu, & Kalra, 2011; Prelec, Wernerfelt, & Zettelmeyer, 1997). That being said, the present study deliberately develops an innovative experiment design and uses non-classical measurement paradigms in order to properly test the current models' assumptions.

No convincing evidence is found for an association between self-reported levels of knowledge and a tendency to choose extreme products in an ordered menu. Knowledge per se seems unable to reduce or eliminate the compromise effect. The results thus fail to provide circumstantial evidence to back up the models' predictions about the alleviating role of consumption experience in the compromise effect. In addition, this finding clearly contradicts past results, which suggest that consumers who are more knowledgeable to (or more familiar with) a certain product category are less likely to choose the compromise option in that category (Sheng, Parker, & Nakamoto, 2005) and are more likely to express consistent preferences about products across choice tasks (Kramer, 2007). One possible explanation of the conflict in empirical results is that Sheng, Parker, and Nakamoto's experiment (2005) adopts a forced-choice design, such that all choice tasks are mandatory. Therefore, their finding may occur because knowledge can effectively steer participants who are uncertain about their preferences, products, etc. away from making a compromise choice through reducing uncertainty. This surmise is partially supported by evidence suggesting that the compromise effect is stronger in presence of uncertainty (Chernev, 2004; Simonson, 1989) and in forced-choice decision tasks (Dhar & Simonson, 2003).

Furthermore, the results show that participants tend to believe in the rational design of product lines[clxviii] and estimate a bell-shaped distribution of market demand over market alternatives. Therefore, rational participants may not directly match their

---

[clxviii] That is, the menu is believed to convey information about the absolute location of other consumers' preferences.

estimated relative standing in taste distribution to a product's relative rank in the market to make a purchasing decision, since the distribution of market demand over available products is perceived as non-uniform. Therefore, choice predictions made from the rank order decision rule require consideration of consumers' estimates of the population distribution of market choices. This suggestion in turn demonstrates that the compromise tendency in choice may not be purely driven by the false consensus effect in preferences. The bell-shaped estimated choice distributions also play a role.

The conjecture that the direct matching strategy is inferior to the comprehensive approach as an account of actual choices did not receive empirical support. The direct matching method predicts purchase choice better than the comprehensive approach does. Seemingly, participants are more likely to directly match their relative standings in the distribution of population's tastes to relative ranks of market alternatives, without considering market demand over products. However, as assumed in the proposed models, the final choice is often not solely determined by the rank order decision rule. The observed contradiction may merely reflect the fact that there exist other drivers of choice. Again, although there is a considerable amount of evidence in favour of the rank order decision rule in the literature (e.g., Burson, 2007; Gershoff & Burson, 2011; Prelec, Wernerfelt, & Zettelmeyer, 1997), the present experiment study is the first one that examines the relative success of the two versions of the rank order decision rule.

Analysis of choice and preferences$_E$ of products empirically confirmed that preferences$_E$ are consistent with choice. Firstly, the average preference ratings of chosen products are higher than those of other products for all product categories. In addition, a substantial fraction of participants rated the chosen products higher than any other choice alternative. Together, results appear to suggest that choice and preferences$_E$ mirror each other. However, the finding that chosen products usually receive higher preference ratings does not necessarily mean that participants choose product with higher preferences$_E$. Since preferences$_E$ were measured after choice, the possibility that participants modified their preferences$_E$ to align with past choice cannot be ruled out. This possibility is consistent with much experimental research in psychology (e.g., Sharot, Fleming, Yu, Koster, & Dolan, 2012; Sharot, Velasquez, & Dolan, 2010; Johansson, Hall, Sikström, & Olsson, 2005) and biology (e.g., Sharot, de Martino, & Dolan, 2009). Therefore, the interpretation of the results about the connection between preferences_E and choice needs to be cautious.

The study has a number of limitations. First, as noted before, the experiment cannot directly test the predictions of the proposed models, owing to technological limitations such as an inability to provide a real product to participants for learning and the difficulty in adopting a repeated choice paradigm (where participants need to repeatedly perform same choice tasks). Secondly, the nature of an online experiment makes it hard to ensure that participants are attending during the experiment, although I have tried to minimise this problem by checking the completion time. Thirdly, without knowing participants' preferences_U, it is impossible to directly examine whether the rank order decision rule can successfully guide people to make an optimal decision as well as how inaccurate estimations about one's own relative position in the population distribution (with respect to tasks) affect choice optimality. Last, the compromise effect found in the experiment only suggests that consumers tend to prefer products placed close to the middle in a menu. This observation, however, is not enough in itself to show whether the compromise effect results from a deviation of choice from preferences_U (or preferences_I).

Chapter 5: Conclusion

Individuals often depart from the narrow sense of rationality assumed in mainstream economics. One example is seen in people's susceptibility to context effects, such as the classic compromise effect investigated in this thesis. Such findings apparently violate the predictions of standard economics, which assumes that rational agents will make choices directly based on their pre-existing underlying preferences that satisfy a set of standard axioms, and hence that choices should be context-independent.

An important issue is the extent to which the compromise effect survives experience. In the case of consumer choice, for example, there are implications for whether or not firms will be able to exploit consumers' susceptibility to context effects. If the compromise effect disappears with sufficient learning in the marketplace, firms will not be able to exploit experienced consumers in the market equilibrium (also they might be able to exploit some fraction of consumers who are an experienced). If on the other hand it survives learning and market experience, there may be an opportunity that firms can take advantage of, which likely resulting in consumer detriment.

A key limitation of most existing model of context effects, the compromise effect in particular, is that they fail to consider the effects of experience and learning in a sufficiently explicit way, and hence do not shed sufficient light on the above issues. This thesis therefore aimed to address three key questions regarding contextual influences on choice. (1) Can people with context-independent underlying preferences nevertheless be susceptible to influence by market/choice context and, if so, why? (2) Under what conditions will context effects (not) persist in market equilibrium even when consumers have extensive market experience? Moreover, is there any factor other than preferences$_U$ that affects final purchasing choices and may potentially contribute to the compromise behaviour? (3) What is the potential impact on consumers' well-being if a profit-maximising firm reacts optimally to consumer bounded rationality?

To address these questions, two models were developed and an experiment was conducted. The first model explored the possibility that the compromise effect arises not from attribute trade-offs, uncertainty about market options or unstable preferences, but instead from biased preference learning in the market due to false consensus effects. Developing Wernerfelt (1995), the model introduced a new preference type, preferences$_N$ (norm-based preferences), to represent the idea that consumers treat menus of market options as informative and tend to leverage them to assist decision making. Simulations of this model demonstrate that the compromise effects may exist in market equilibrium (after learning) when consumers choose rationally based on their learned preferences. The model simulations also show that it can be profitable for a monopolist to distort menus. However, this conclusion was highly dependent on choice of parameter values and market contexts.

A second model addressed a number of limitations of the first model. In particular, (a) the first model predicted implausible patterns of consumer choices as the attribute values of available products move towards the upper end of the attribute space; (b) the effect of market-based inferences was rather weak in the first model, and (c) the first model adopted an implausible method of preference learning which, once complete, did not allow any further change.

Model 2 retained the assumptions about preferences$_U$ and preferences$_N$ with false consensus, while employing a classical explore-exploit framework. The approach assumes that at any given choice point, false-consensus consumers will either exploit (on the basis of their preference$_I$, which is formed by their past choices) or explore (on the basis of preferences$_N$). An additional feature was that item choice will lead to subsequent inhibition of market options, making them less likely to be chosen on subsequent occasions. Computational exploration of this second model reproduced the compromise effect and showed that firms may have an opportunity to take advantage of biased consumers by distorting market of options available to consumers.

Finally, an empirical study offered an initial exploration of the validity of the basic assumptions of the models. Results generally supported the assumption that people tend to believe that they are nearer to the average of the population in terms of tastes than they actually are. The sizes of estimation errors were, however, not significantly associated with self-reported levels of knowledge. Secondly, a tendency for participants to choose products placed around the middle of the ordered set was observed for all product categories, indicating presence of a compromise effect. Thirdly, a series of tests of the assumption of the rank order decision rule indicated that participants normally expect a rational design of product lines, but assume a bell-shaped distribution of market demand for products. However, experimental results showed that the direct matching strategy better predicted choices for all product categories.

5.1 Conclusion

This thesis adds to the existing literature by offering an account of how the compromise effect may reflect rational responses to information provided by market contexts. In relation to the three research questions outlined above, the following conclusions are drawn. First, regarding the question of why people with context-independent preferences$_{U}$ may still exhibit the compromise behaviour, it is concluded that the conscious inaccessibility of absolute values of the preference$_{U}$ that motivates people to learn their true preferences in the market through consuming market offerings may lead them to believe that the middle option is optimal. This is because individuals who treat the product line as informative and tend to use it to facilitate decision making are likely to generate biased estimates of their relative standing in the population's taste distribution due to the false consensus effect. This conclusion was established in Simulation 1 (Section 2.3.2.2), Simulation 2 (Section 2.3.3.1 and 2.3.3.2) and Simulation 3 (Section 2.3.4.2) in Chapter 3 as well as Simulation 1 (Section 3.4.2.3), Simulation 2 (Section 3.4.3.2) and Simulation 4 (Section 3.4.5.1) in Chapter 4. Specifically, in both models, all simulation results show that decision makers who do not use the rank-order decision rule at all act if they are fully aware of their preference$_{U}$ and can make decisions solely based on it. Conversely, consumers who employ a market inference strategy, deriving from flawed estimates of their own relative position in the population distribution, display a strong tendency to choose the compromise option, despite having a context-independent preference$_{U}$. Accordingly, the present thesis argue that context-dependent behaviour does not necessarily stem from a lack of innate, stable preferences$_{U}$ and it thus cannot be used as an evidence of preference instability or irrationality.

The second question that concerns conditions for the compromise effect to persist or disappear in market equilibrium where consumers have extensive market experience can be answered by the simulation results in the present thesis. Briefly speaking, results obtained in Simulation 2 (Section 2.3.3.1 and 2.3.3.2) and Simulation 3 (Section 2.3.4.2) in Chapter 2 and Simulation 2 (Section 3.4.3.3) and Simulation 4 (Section 3.4.5.3, 3.4.5.4 and 3.4.5.5) in Chapter 3 clearly lead to the conclusion that experience *per se* cannot eliminate the compromise effect as long as preference learning in the market is biased (due to the false consensus effect). The proposed models therefore contrast with existing theories that predict that the compromise effect will inevitably disappear when uncertainty is resolved. Yet, phenomena such as the observed upward trend in food portion size and intake (see Chapter 1) may reflect the possibility that context effects, including the compromise effect, may persist in a market that involves experienced consumers.

Further, results presented in Chapter 2 show that the compromise effect will no longer appear when market contexts are too unfavourable with respect to decision makers' preferences$_U$. This is because unpleasant signals sent from the preference$_U$ become more salient when market alternatives get more unsatisfactory. Simulation results in Chapter 3, however, suggest that the changes in the intensity of the compromise effect due to contextual changes depend largely on the measurement approaches. Therefore, there is no certain pattern about changes of the compromise effect predicted in Model 2. That being said, all simulation results in Chapter 3 show that the compromise effect will never disappear because in the second model, consumers' unpleasant memories about bad products decays over time, which make them possibly choose the undesirable compromise option when exploring, especially for those who have a false consensus bias.

In addition, the extended second question which asks about any factor other than preferences$_U$ that influences final choices and may potentially yield the compromise behaviour is also answered in the thesis. As noted above, all simulation results in Chapters 2 and 3 clearly demonstrate that social norms may lead to the compromise effect through biasing market inferences. In the two proposed models, social norms is the key factor that purges the compromise effect in equilibrium under some (or all) market contexts.

Last, the question concerning the monopolist's profit-maximising strategy in setting the product line, and the effect of this strategic reaction on consumers' welfare in equilibrium is also solved. The thesis concludes that when the firm can benefit from selling products with larger attribute values, there is an incentive to exploit the equilibrium compromise effect and incur the detriment of consumers. Conversely, in the case where prices and marginal costs are constant across all feasible products, the monopolist will not engage in context distortion since all products bring equal profits. This conclusion is supported by results in Simulation 3 (Section 2.3.4.4 and 2.3.4.5) in Chapter 2 and in Simulation 1 (Section 3.4.2.5, 3.4.2.6 and 3.4.2.7), Simulation 3 (Section 3.4.4.1, 3.4.4.2 and 3.4.4.3), and Simulation 4 (Section 3.4.5.6, 3.4.5.7 and 3.4.5.8) in Chapter 3. This result is consistent with Kamenica (2008) and Sharp, Staelin and Huber's (2008) theoretical prediction, which suggests that a profit-maximising firm may introduce a premium loss leader to the market, aiming to promote sales of more profitable products. Kamenica's (2008) finding, however, only holds when a sufficient amount of inexperienced consumers exists in the market. Our results complement Kamenica's (2008) study by showing that context distortion can occur in the market full of experienced consumers.

5.2 Limitations, directions, and relation to previous research

The experiment differs from previous studies on the compromise effect in a number of ways. Firstly, in contrast to many empirical and theoretical studies of the compromise effect that use more than one product attribute (e.g., Pocheptsova, Amir, Dhar, & Baumeister 2009; Simonson & Tversky, 1992), products within a category only differed in one attribute dimension in the present studies. Thus, there was no manipulation or exploration of attribute trade-offs. Secondly, throughout the experiment, numeric attribute representations of products were not presented to participants. Instead, only information about the ordinal position of products was shown. This design was employed to avoid unexpected influences of numerical values on participants' evaluations and choices.

Thirdly, the present experiment did not manipulate the effects of contextual change on choice of options, i.e., it did not use a between-subject design to compare the choice share of the same product in different sizes of choice sets. This is because one of the arguments made by the proposed models is that the compromise effect may still be observed when contextual changes do not involve a change in the size of a choice set[clxix] . Hence, in order to test the models' predictions about consumer choice with present experiment data, sizes of choice sets are best kept constant, whether context is manipulated or not. In addition, since attribute values of products were hidden from participants, comparing choice shares gained from two seemingly identical sets is senseless, unless the experiment adopts a within-subject design and participants are instructed that one set of products is, on average, larger in attribute magnitudes than another. Last, unlike most existing experiments on the compromise effect (e.g., Prelec, Wernerfelt, & Zettelmeyer, 1997; Sheng, Parker, & Nakamoto, 2005), participants in the present experiment were given a no-choice option for all questions about purchase decisions. This unforced choice paradigm was used to avoid the possibility that the compromise effect is driven by a feeling of difficulty in choice tasks.

Future work will be needed to address a number of limitations of the current research on modelling the compromise effect. First, a more comprehensive model will need to incorporate the idea that exploration-based choices will likely become relatively less frequent as consumers become increasingly experienced within a domain. The extent to which reduction in exploration maximizes utility will of course depend on the frequency with which choosers' preferences on the one hand, and the choice environment on the other, change over time.

Secondly, and relatedly, incentives on firms will depend greatly on the proportion of inexperienced and experienced consumers within the market place. Further theoretical explanations will be necessary to get a full understanding of the relevant contingencies.

Thirdly, a major limitation of the current work is that it assumes the existence of a single firm which has a monopoly advantage. It is of practical importance to understand

---

[clxix] As discussed in Chapter 1, many theoretical accounts of the compromise effect (e.g., Guo, 2016; Kamenica, 2008) only work in a situation where the size of a choice set was enlarged after a contextual change. Therefore, they may fail to explain or predict a possible market scenario where the compromise effect arises when the size of menu remains the same after change.

the conditions under which the exploitation potential that the present work identifies can be removed or reduced by market competition and/or regulation.

References

Afriat, S. N. (1967). The construction of utility functions from expenditure data. *International Economic Review, 8*(1), 67-77.

Albar, S. A., Alwan, N. A., Evans, C. E., & Cade, J. E. (2014). Is there an association between food portion size and BMI among British adolescents? *British Journal of Nutrition*, *112*(5), 841-851.

Aldrovandi, S., Brown, G. D. A., & Wood, A. M. (2015). Social norms and rank-based nudging: Changing willingness to pay for healthy food. *Journal of Experimental Psychology: Applied*, *21*(3), 242-252.

Alekseev, A., Charness, G., & Gneezy, U. (2017). Experimental methods: When and why contextual instructions are important. *Journal of Economic Behavior & Organization*, *134*, 48-59.

Amir, O., & Levav, J. (2008). Choice construction versus preference construction: The instability of preferences learned in context. *Journal of Marketing Research*, *45*(2), 145-158.

Ariely, D., Loewenstein, G., & Prelec, D. (2003). "Coherent arbitrariness": Stable demand curves without stable preferences. *The Quarterly Journal of Economics*, *118*(1), 73-106.

Arrow, K. J. (1951). *Social choice and individual values*. New York, NY: John Wiley & Sons, Inc..

Auer, P. (2002). Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, *3*, 397-422.

Becker, G. M., DeGroot, M. H., & Marschak, J. (1963). Probabilities of choices among very similar objects: An experiment to decide between two models. *Behavioral Science*, *8*(4), 306-311.

Bem, D. J. (1967). Self-perception: An alternative interpretation of cognitive dissonance phenomena. *Psychological Review*, *74*(3), 183-200.

Bem, D. J. (1972). Self-perception theory. *Advances in Experimental Social Psychology*, *6*(1), 1-62.

Bennett, P. J., & Pratt, J. (2001). The spatial distribution of inhibition of return. *Psychological Science*, *12*(1), 76-80.

Benton, D. (2015). Portion size: What we know and what we need to know. *Critical Reviews in Food Science and Nutrition*, *55*(7), 988-1004.

Bettman, J. R., Luce, M. F., & Payne, J. W. (1998). Constructive consumer choice processes. *Journal of Consumer Research*, *25*(3), 187-217.

Bettman, J. R., Luce, M. F., & Payne, J. W. (2008). Preference construction and preference stability: Putting the pillow to rest. *Journal of Consumer Psychology*, *18*(3), 170-174.

Binmore, K. (2008). *Rational Decisions*. Oxfordshire, the UK: Princeton University Press.

Birmingham, E., Visser, T. A., Snyder, J. J., & Kingstone, A. (2007). Inhibition of return: Unraveling a paradox. *Psychonomic Bulletin & Review*, *14*(5), 957-963.

Blanco, N. J., Love, B. C., Cooper, J. A., McGeary, J. E., Knopik, V. S., & Maddox, W. T. (2015). A frontal dopamine system for reflective exploratory behavior. *Neurobiology of Learning and Memory*, 123, 84-91.

Bordalo, P., Gennaioli, N., & Shleifer, A. (2013). Salience and consumer choice. *Journal of Political Economy*, *121*(5), 803-843.

Bordalo, P., Gennaioli, N., & Shleifer, A. (2020). Memory, attention, and choice. *The Quarterly Journal of Economics*, *135*(3), 1399-1442.

Bornstein, A. M., & Norman, K. A. (2017). Reinstated episodic context guides sampling-based decisions for reward. *Nature Neuroscience*, *20*(7), 997-1003.

Boyce, C. J., Brown, G. D. A., & Moore, S. C. (2010). Money and happiness: Rank of income, not income, affects life satisfaction. *Psychological Science*, *21*(4), 471-475.

Braga, J., Humphrey, S. J., & Starmer, C. (2009). Market experience eliminates some anomalies—and creates new ones. *European Economic Review*, *53*(4), 401-416.

Brouwer, R. (2012). Constructed preference stability: a test–retest. *Journal of Environmental Economics and Policy*, *1*(1), 70-84.

Brouwer, R., Dekker, T., Rolfe, J., & Windle, J. (2010). Choice certainty and consistency in repeated choice experiments. *Environmental and Resource Economics*, 46(1), 93-109.

Brown, G. D. A., & Walasek, L. (2018). *Relative rank theory: Preferences, choices, and the inaccessibility of value*. Unpublished manuscript.

Brown, T. C., Kingsley, D., Peterson, G. L., Flores, N. E., Clarke, A., & Birjulin, A. (2008). Reliability of individual valuations of public and private goods: Choice consistency, response time, and preference refinement. *Journal of Public Economics*, *92*(7), 1595-1606.

Bruine de Bruin, W., Galesic, M., Parker, A. M., & Vardavas, R. (2020). The role of social circle perceptions in "false consensus" about population statistics: Evidence from a national flu survey. *Medical Decision Making*, *40*(2), 235-241.

Burson, K. A. (2007). Consumer-product skill matching: The effects of difficulty on relative self-assessment and choice. *Journal of Consumer Research*, *34*(1), 104-110.

Burson, K. A., & Gershoff, A. D. (2015). Marketing actions that influence estimates of others also shape identity. *Journal of Consumer Psychology*, *25*(3), 495-503.

Burson, K. A., Larrick, R. P., & Klayman, J. (2006). Skilled or unskilled, but still unaware of it: How perceptions of difficulty drive miscalibration in relative comparisons. *Journal of Personality and Social Psychology*, *90*(1), 60-77.

Byrne, B. M. (2013). *Structural equation modeling with Mplus: Basic concepts, applications, and programming*. Oxford, the U.K.: Routledge.

Carlson, K. A., & Bond, S. D. (2006). Improving preference assessment: Limiting the effect of context through pre-exposure to attribute levels. *Management Science*, *52*(3), 410-421.

Carroll, R., & Vallen, B. (2014). Compromise and attraction effects in food choice. *International Journal of Consumer Studies*, *38*(6), 636-641.

Cavanna, A. E., Shah, S., Eddy, C. M., Williams, A., & Rickards, H. (2011). Consciousness: A neurological perspective. *Behavioural Neurology*, *24*(1), 107-116.

Cesarini, D., Lichtenstein, P., Johannesson, M., & Wallace, B. (2009). Heritability of overconfidence. *Journal of the European Economic Association*, *7*(2-3), 617-627.

Chandon, P., & Wansink, B. (2011). Is food marketing making us fat? A multi-disciplinary review. *Foundations and Trends® in Marketing*, *5*(3), 113-196.

Chandon, P., & Wansink, B. (2012). Does food marketing need to make us fat? A review and solutions. *Nutrition Reviews*, *70*(10), 571-593.

Chartrand, T. L., & Bargh, J. A. (1999). The chameleon effect: The perception–behavior link and social interaction. *Journal of Personality and Social Psychology*, *76*(6), 893-910.

Chernev, A. (2004). Extremeness aversion and attribute-balance effects in choice. *Journal of Consumer Research*, *31*(2), 249-263.

Christie, J., Hilchey, M. D., & Klein, R. M. (2013). Inhibition of return is at the midpoint of simultaneous cues. *Attention, Perception, & Psychophysics*, *75*(8), 1610-1618.

Chuang, S. C., Cheng, Y. H., Chang, C. J., & Chiang, Y. T. (2013). The impact of self-confidence on the compromise effect. *International Journal of Psychology*, *48*(4), 660-675.

Chuang, S. C., Kao, D. T., Cheng, Y. H., & Chou, C. A. (2012). The effect of incomplete information on the compromise effect. *Judgment and Decision Making*, *7*(2), 196-206.

Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1481), 933-942.

Coupey, E., Irwin, J. R., & Payne, J. W. (1998). Product category familiarity and preference construction. *Journal of Consumer Research*, *24*(4), 459-468.

Cox, J. C., & Grether, D. M. (1996). The preference reversal phenomenon: Response mode, markets and incentives. *Economic Theory*, *7*(3), 381-405.

Cronqvist, H., & Siegel, S. (2014). The genetics of investment biases. *Journal of Financial Economics*, *113*(2), 215-234.

Currarini, S., & Mengel, F. (2016). Identity, homophily and in-group bias. *European Economic Review*, *90*, 40-55.

Dawtry, R. J., Sutton, R. M., & Sibley, C. G. (2015). Why wealthier people think people are wealthier, and why it matters: From social sampling to attitudes to redistribution. *Psychological Science*, *26*(9), 1389-1400.

Dayan, P., & Berridge, K. C. (2014). Model-based and model-free Pavlovian reward learning: revaluation, revision, and revelation. *Cognitive, Affective, & Behavioral Neuroscience*, *14*(2), 473-492.

de Bekker-Grob, E. W., & Chorus, C. G. (2013). Random regret-based discrete-choice modelling: An application to healthcare. *PharmacoEconomics*, *31*(7), 623-634.

Denrell, J. (2005). Why most people disapprove of me: experience sampling in impression formation. *Psychological Review, 112*(4), 951-978.

Deutsch, M., & Gerard, H. B. (1955). A study of normative and informational social influences upon individual judgment. *The Journal of Abnormal and Social Psychology, 51*(3), 629-636.

Dhar, R. (1996). The effect of decision strategy on deciding to defer choice. *Journal of Behavioral Decision Making*, *9*(4), 265-281.

Dhar, R., & Gorlin, M. (2013). A dual-system framework to understand preference construction processes in choice. *Journal of Consumer Psychology*, *23*(4), 528-542.

Dhar, R., & Sherman, S. J. (1996). The effect of common and unique features in consumer choice. *Journal of Consumer Research*, 23(3), 193-203.

Dhar, R., & Simonson, I. (2003). The effect of forced choice on choice. *Journal of Marketing Research*, *40*(2), 146-160.

Diliberti, N., Bordi, P. L., Conklin, M. T., Roe, L. S., & Rolls, B. J. (2004). Increased portion size leads to increased energy intake in a restaurant meal. *Obesity Research*, *12*(3), 562-568.

Dobson, P. W., & Gerstner, E. (2010). For a few cents more: Why supersize unhealthy food?. *Marketing Science*, *29*(4), 770-778.

Drolet, A., Luce, M. F., & Simonson, I. (2008). When does choice reveal preference? Moderators of heuristic versus goal-based choice. *Journal of Consumer Research*, *36*(1), 137-147.

Duncan, R. B. (1972). Characteristics of organizational environments and perceived environmental uncertainty. *Administrative Science Quarterly*, *17*(3), 313-327.

Durante, K. M., & Griskevicius, V. (2016). Evolution and consumer behavior. *Current Opinion in Psychology*, *10*, 27-32.

English, L., Lasschuijt, M., & Keller, K. L. (2015). Mechanisms of the portion size effect. What is known and where do we go from here? *Appetite*, *88*, 39-49.

Erev, I., Ert, E., Plonsky, O., Cohen, D., & Cohen, O. (2017). From anomalies to forecasts: Toward a descriptive model of decisions under risk, under ambiguity, and from experience. *Psychological Review*, *124*(4), 369-409.

Ert, E., & Lejarraga, T. (2018). The effect of experience on context-dependent decisions. *Journal of Behavioral Decision Making*, *31*(4), 535-546.

Feng, S. F., Wang, S., Zarnescu, S., & Wilson, R. C. (2021). The dynamics of explore–exploit decisions reveal a signal-to-noise mechanism for random exploration. *Scientific Reports*, *11*(1), 1-15.

Fiedler, K. (2000). Beware of samples! A cognitive-ecological sampling approach to judgment biases. *Psychological Review*, *107*(4), 659-676.

Flood, J. E., Roe, L. S., & Rolls, B. J. (2006). The effect of increased beverage portion size on energy intake at a meal. *Journal of the American Dietetic Association*, *106*(12), 1984-1990.

Frederick, S., Lee, L., & Baskin, E. (2014). The limits of attraction. *Journal of Marketing Research*, *51*(4), 487-507.

Freund, P. A., & Kasten, N. (2012). How smart do you think you are? A meta-analysis on the validity of self-estimates of cognitive ability. *Psychological Bulletin*, *138*(2), 296-321.

Friedman, D., Isaac, R. M., James, D., & Sunder, S. (2014). *Risky curves: On the empirical failure of expected utility*. London: Routledge.

Fudenberg, D., & Kreps, D. M. (1993). Learning mixed equilibria. *Games and Economic Behavior*, *5*(3), 320-367.

Funk, C. L., Smith, K. B., Alford, J. R., Hibbing, M. V., Eaton, N. R., Krueger, R. F., Eaves, L. J., & Hibbing, J. R. (2013). Genetic and environmental transmission of political orientations. *Political Psychology*, *34*(6), 805-819.

Galesic, M., Olsson, H., & Rieskamp, J. (2012). Social sampling explains apparent biases in judgments of social environments. *Psychological Science*, *23*(12), 1515-1523.

Galesic, M., Olsson, H., & Rieskamp, J. (2018). A sampling model of social judgment. *Psychological Review*, *125*(3), 363-390.

George, D., & Mallery, P. (2019). *IBM SPSS statistics 26 step by step: A simple guide and reference*. Oxford, the U.K.: Routledge.

Gerasimou, G. (2017). Indecisiveness, undesirability and overload revealed through rational choice deferral. *The Quarterly Journal of Economics*, *125*(4), 1399-1433.

Gershman, S. J. (2018). The successor representation: its computational logic and neural substrates. *Journal of Neuroscience*, *38*(33), 7193-7200.

Gershoff, A. D., & Burson, K. A. (2011). Knowing where they stand: The role of inferred distributions of others in misestimates of relative standing. *Journal of Consumer Research*, *38*(3), 407-419.

Gershoff, A. D., Mukherjee, A., & Mukhopadhyay, A. (2007). What's not to like? Preference asymmetry in the false consensus effect. *Journal of Consumer Research*, *35*(1), 119-125.

Grether, D. M., & Plott, C. R. (1979). Economic theory of choice and the preference reversal phenomenon. *American Economic Review*, *69*(4), 623-638.

Guo, L. (2016). Contextual deliberation and preference construction. *Management Science*, *62*(10), 2977-2993.

Hadar, L., Danziger, S., & Hertwig, R. (2018). The attraction effect in experience-based decisions. *Journal of Behavioral Decision Making*, *31*(3), 461-468.

Hair, J. F., Black, W. C., Babin, B. J., & Anderson, R. E. (2013). *Multivariate data analysis: Pearson new international edition*. London: Pearson Education, Limited

Hall, L., Johansson, P., Tärning, B., Sikström, S., & Deutgen, T. (2010). Magic at the marketplace: Choice blindness for the taste of jam and the smell of tea. *Cognition*, *117*(1), 54-61.

Hamilton, R. W., Ratner, R. K., & Thompson, D. V. (2010). Outpacing others: When consumers value products based on relative usage frequency. *Journal of Consumer Research*, *37*(6), 1079-1094.

Hausman, D. M. (1992). *Essays on philosophy and economic methodology*. Cambridge: Cambridge University Press.

Helzer, E. G., & Dunning, D. (2012). Why and when peer prediction is superior to self-prediction: The weight given to future aspiration versus past achievement. *Journal of Personality and Social Psychology, 103*(1), 38-53.

Herman, C. P., Polivy, J., Pliner, P., & Vartanian, L. R. (2015). Mechanisms underlying the portion-size effect. *Physiology & Behavior, 144*, 129-136.

Hertwig, R., & Pleskac, T. J. (2010). Decisions from experience: Why small samples?. *Cognition*, *115*(2), 225-237.

Hertwig, R., Pachur, T., & Kurzenhäuser, S. (2005). Judgments of risk frequencies: Tests of possible cognitive mechanisms. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*(4), 621.

Higgs, S. (2015). Social norms and their influence on eating behaviours. *Appetite*, *86*, 38-44.

Hoch, S. J. (1988). Who do we know: Predicting the interests and opinions of the American consumer. *Journal of Consumer Research*, *15*(3), 315-324.

Hoeffler, S., & Ariely, D. (1999). Constructing stable preferences: A look into dimensions of experience and their impact on preference stability. *Journal of Consumer Psychology*, *8*(2), 113-139.

Hoeffler, S., Ariely, D., West, P., & Duclos, R. (2013). Preference exploration and learning: The role of intensiveness and extensiveness of experience. *Journal of Consumer Psychology*, *23*(3), 330-340.

Houthakker, H. S. (1950). Revealed preference and the utility function. *Economica*, *17*(66), 159-174.

Howes, A., Warren, P. A., Farmer, G., El-Deredy, W., & Lewis, R. L. (2016). Why contextual preference reversals maximize expected value. *Psychological Review*, *123*(4), 368-391.

Hsee, C. K., Yang, Y., Li, N., & Shen, L. (2009). Wealth, warmth, and well-being: Whether happiness is relative or absolute depends on whether it is about money, acquisition, or consumption. *Journal of Marketing Research*, *46*(3), 396-409.

Hu, C., Domenech, P., & Pessiglione, M. (2020). Order matters: How covert value updating during sequential option sampling shapes economic preference. *PLoS Computational Biology*, *16*(8), 1-23.

Huber, J., Payne, J. W., & Puto, C. P. (1982). Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of Consumer Research*, *9*(1), 90-98.

Huber, J., Payne, J. W., & Puto, C. P. (2014). Let's be honest about the attraction effect. *Journal of Marketing Research*, *51*(4), 520-525.

Hur, Y. M., & Rushton, J. P. (2007). Genetic and environmental contributions to prosocial behaviour in 2-to 9-year-old South Korean twins. *Biology Letters*, *3*(6), 664-666.

Hutchinson, J. W., Kamakura, W. A., & Lynch Jr, J. G. (2000). Unobserved heterogeneity as an alternative explanation for "reversal" effects in behavioral research. *Journal of Consumer Research*, *27*(3), 324-344.

Isoni, A., Brooks, P., Loomes, G., & Sugden, R. (2016). Do markets reveal preferences or shape them? *Journal of Economic Behavior & Organization*, *122*, 1-16.

Izuma, K., & Murayama, K. (2013). Choice-induced preference change in the free-choice paradigm: A critical methodological review. *Frontiers in Psychology*, *4*(41), 1-12.

Jackson, J. J., Connolly, J. J., Garrison, S. M., Leveille, M. M., & Connolly, S. L. (2015). Your friends know how long you will live: A 75-year study of peer-rated personality traits. *Psychological Science*, *26*(3), 335-340.

Jehiel, P., & Steiner, J. (2020). Selective sampling with information-storage constraints. *The Economic Journal*, *130*(630), 1753-1781.

Jensen, N. E. (1967). An introduction to Bernoullian utility theory: I. Utility functions. *The Swedish Journal of Economics*, *69*(3), 163-183.

Johansson, P., Hall, L., Sikström, S., & Olsson, A. (2005). Failure to detect mismatches between intention and outcome in a simple decision task. *Science*, *310*(5745), 116-119.

Judd, C. M., Ryan, C. S., & Park, B. (1991). Accuracy in the judgment of in-group and out-group variability. *Journal of Personality and Social Psychology*, *61*(3), 366-379.

Juslin, P., Winman, A., & Hansson, P. (2007). The naïve intuitive statistician: A naïve sampling model of intuitive confidence intervals. *Psychological Review*, *114*(3), 678-703.

Kahn, B. E., Kalwani, M. U., & Morrison, D. G. (1986). Measuring variety-seeking and reinforcement behaviors using panel data. *Journal of Marketing Research*, *23*(2), 89-100.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263-291.

Kakade, S., & Dayan, P. (2002). Dopamine: Generalization and bonuses. *Neural Networks*, *15*(4-6), 549-559.

Kamenica, E. (2008). Contextual inference in markets: On the informational content of product lines. *American Economic Review*, *98*(5), 2127-2149.

Keeling, S. (2018). Confabulation and rational obligations for self-knowledge. *Philosophical Psychology*, *31*(8), 1215-1238.

Kitts, J. A. (2003). Egocentric bias or information management? Selective disclosure and the social roots of norm misperception. *Social Psychology Quarterly*, *66*(3), 222-237.

Kivetz, R., Netzer, O., & Srinivasan, V. (2004). Alternative models for capturing the compromise effect. *Journal of Marketing Research*, *41*(3), 237-257.

Khan, U., Zhu, M., & Kalra, A. (2011). When trade-offs matter: The effect of choice construal on context effects. *Journal of Marketing Research*, *48*(1), 62-71.

Klein, R. M. (2000). Inhibition of return. *Trends in Cognitive Sciences*, *4*(4), 138-147.

Klein, R. M., & MacInnes, W. J. (1999). Inhibition of return is a foraging facilitator in visual search. *Psychological Science*, *10*(4), 346-352.

Kling, S. M., Roe, L. S., Keller, K. L., & Rolls, B. J. (2016). Double trouble: Portion size and energy density combine to increase preschool children's lunch intake. *Physiology & Behavior*, *162*, 18-26.

Klucharev, V., Hytönen, K., Rijpkema, M., Smidts, A., & Fernández, G. (2009). Reinforcement learning signal predicts social conformity. *Neuron*, *61*(1), 140-151.

Knetsch, J. L., Tang, F. F., & Thaler, R. H. (2001). The endowment effect and repeated market trials: Is the Vickrey auction demand revealing? *Experimental Economics*, *4*(3), 257-269.

Krajbich, I., Armel, C., & Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience*, *13*(10), 1292-1298.

Kramer, T. (2007). The effect of measurement task transparency on preference construction and evaluations of personalized recommendations. *Journal of Marketing Research*, *44*(2), 224-233.

Krueger, J., & Mueller, R. A. (2002). Unskilled, unaware, or both? The better-than-average heuristic and statistical regression predict errors in estimates of own performance. *Journal of Personality and Social Psychology*, *82*(2), 180-188.

Kruger, J. (1999). Lake Wobegon be gone! The" below-average effect" and the egocentric nature of comparative ability judgments. *Journal of Personality and Social Psychology*, 77(2), 221-232.

Latty, T., & Beekman, M. (2011). Irrational decision-making in an amoeboid organism: transitivity and context-dependent preferences. *Proceedings of the Royal Society B: Biological Sciences*, *278*(1703), 307-312.

Lea, A. M., & Ryan, M. J. (2015). Irrationality in mate choice revealed by túngara frogs. *Science*, *349*(6251), 964-966.

Leong, W., & Hensher, D. A. (2014). Relative advantage maximisation as a model of context dependence for binary choice data. *Journal of Choice Modelling*, *11*, 30-42.

Levav, J., & Zhu, R. (2009). Seeking freedom through variety. *Journal of Consumer Research*, *36*(4), 600-610.

Levin, I. P., Gaeth, G. J., Schreiber, J., & Lauriola, M. (2002). A new look at framing effects: Distribution of effect sizes, individual differences, and independence of types of effects. *Organizational Behavior And Human Decision Processes*, *88*(1), 411-429.

Liang, S., Gemming, L., Wellard-Cole, L., & Rangan, A. (2019). Comparison between serving sizes of cakes and muffins sold in Australian supermarkets and coffee shop chains. *Nutrition & Dietetics*, *76*(3), 284-289.

Lichtenstein, S., & Slovic, P. (Eds.). (2006). *The construction of preference*. Cambridge: Cambridge University Press.

Lichters, M., Brunnlieb, C., Nave, G., Sarstedt, M., & Vogt, B. (2016). The influence of serotonin deficiency on choice deferral and the compromise effect. *Journal of Marketing Research*, *53*(2), 183-198.

Lichters, M., Sarstedt, M., & Vogt, B. (2015). On the practical relevance of the attraction effect: A cautionary note and guidelines for context effect experiments. *AMS Review*, *5*(1-2), 1-19.

List, J. A. (2003). Does market experience eliminate market anomalies?. *Quarterly Journal of Economics*, *118*(1), 41–71.

List, J. A. (2004). Neoclassical theory versus prospect theory: Evidence from the marketplace. *Econometrica*, *72*(2), 615-625.

List, J. A. (2011). Does market experience eliminate market anomalies? The case of exogenous market experience. *American Economic Review*, *101*(3), 313-317.

Livingstone, M. B. E., & Pourshahidi, L. K. (2014). Portion size and obesity. *Advances in Nutrition*, *5*(6), 829-834.

Loomes, G., Starmer, C., & Sugden, R. (2003). Do anomalies disappear in repeated markets? *The Economic Journal*, *113*(486), C153-C166.

Loomes, G., Starmer, C., & Sugden, R. (2010). Preference reversals and disparities between willingness to pay and willingness to accept in repeated markets. *Journal of Economic Psychology*, *31*(3), 374-387.

Luce, R. D. (1959). On the possible psychophysical laws. *Psychological Review*, *66*(2), 81-95.

Luce, R. D. (1977). The choice axiom after twenty years. *Journal of Mathematical Psychology*, *15*(3), 215-233.

Madan, C. R., Ludvig, E. A., & Spetch, M. L. (2014). Remembering the best and worst of times: Memories for extreme outcomes bias risky decisions. *Psychonomic Bulletin & Review*, *21*(3), 629-636.

Maniadis, Z., Tufano, F., & List, J. A. (2014). One swallow doesn't make a summer: New evidence on anchoring effects. *American Economic Review*, *104*(1), 277-90.

Manzini, P., & Mariotti, M. (2014). Stochastic choice and consideration sets. *Econometrica*, *82*(3), 1153-1176.

Marchiori, D., Waroquier, L., & Klein, O. (2011). Smaller food item sizes of snack foods influence reduced portions and caloric intake in young adults. *Journal of the American Dietetic Association*, *111*(5), 727-731.

Mariotti, M. (2008). What kind of preference maximization does the weak axiom of revealed preference characterize?. *Economic Theory*, *35*(2), 403-406.

Markant, D. B., & Gureckis, T. M. (2014). Is it better to select or to receive? Learning via active and passive hypothesis testing. *Journal of Experimental Psychology: General*, *143*(1), 94-122.

Marks, G., & Miller, N. (1987). Ten years of research on the false-consensus effect: An empirical and theoretical review. *Psychological Bulletin*, *102*(1), 72-90.

Marschak, J. (1950). Rational behavior, uncertain prospects, and measurable utility. *Econometrica*, *18*(2) 111-141.

Matthiessen, J., Fagt, S., Biltoft-Jensen, A., Beck, A. M., & Ovesen, L. (2003). Size makes a difference. *Public Health Nutrition*, *6*(1), 65-72.

McFadden, D. (1999). Rationality for Economists?. *Journal of Risk and Uncertainty*, *19*, 73-105.

McFadden, D. (2001). Economic choices. *American Economic Review*, *91*(3), 351-378.

McFerran, B., Dahl, D. W., Fitzsimons, G. J., & Morales, A. C. (2010). Might an overweight waitress make you eat more? How the body type of others is sufficient to alter our food consumption. *Journal of Consumer Psychology*, *20*(2), 146-151.

McPherson, M., Smith-Lovin, L., & Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, *27*(1), 415-444.

Mobbs, D., Hassabis, D., Yu, R., Chu, C., Rushworth, M., Boorman, E., & Dalgleish, T. (2013). Foraging under competition: The neural basis of input-matching in humans. *Journal of Neuroscience*, *33*(23), 9866-9872.

Mochon, D. (2013). Single-option aversion. *Journal of Consumer Research*, *40*(3), 555-566.

Mullen, B., & Hu, L. T. (1988). Social projection as a function of cognitive mechanisms: Two meta-analytic integrations. *British Journal of Social Psychology*, *27*(4), 333-356.

Mullen, B., Atkins, J. L., Champion, D. S., Edwards, C., Hardy, D., Story, J. E., & Vanderklok, M. (1985). The false consensus effect: A meta-analysis of 115 hypothesis tests. *Journal of Experimental Social Psychology*, *21*(3), 262-283.

Mullen, B., Dovidio, J. F., Johnson, C., & Copper, C. (1992). In-group-out-group differences in social projection. *Journal of Experimental Social Psychology*, *28*(5), 422-440.

Müller, H., Kroll, E. B., & Vogt, B. (2012). Do real payments really matter? A re-examination of the compromise effect in hypothetical and binding choice settings. *Marketing Letters*, *23*(1), 73-92.

Natenzon, P. (2019). Random choice and learning. *Journal of Political Economy*, *127*(1), 419-457.

Nestle, M. (2003). Increasing portion sizes in American diets: More calories, more obesity. *Journal of the Academy of Nutrition and Dietetics*, 103(1), 39-40.

Neumann, N., Böckenholt, U., & Sinha, A. (2016). A meta-analysis of extremeness aversion. *Journal of Consumer Psychology*, *26*(2), 193-212.

Niedrich, R. W., Weathers, D., Hill, R. C., & Bell, D. R. (2009). Specifying price judgments with range–frequency theory in models of brand choice. *Journal of Marketing Research*, *46*(5), 693-702.

Nielsen, S. J., & Popkin, B. M. (2003). Patterns and trends in food portion sizes, 1977-1998. *JAMA*, *289*(4), 450-453.

Nisbett, R. E., & Kunda, Z. (1985). Perception of social distributions. *Journal of Personality and Social Psychology*, *48*(2), 297-311.

Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, *84*(3), 231-259.

Noguchi, T., & Stewart, N. (2014). In the attraction, compromise, and similarity effects, alternatives are repeatedly compared in pairs on single dimensions. *Cognition*, *132*(1), 44-56.

O'Boyle Jr, E., & Aguinis, H. (2012). The best and the rest: Revisiting the norm of normality of individual performance. *Personnel Psychology*, *65*(1), 79-119.

Otto, A. R., Knox, W. B., Markman, A. B., & Love, B. C. (2014). Physiological and behavioral signatures of reflective exploratory choice. *Cognitive, Affective, & Behavioral Neuroscience, 14*(4), 1167-1183.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1992). Behavioral decision research: A constructive processing perspective. *Annual Review of Psychology, 43*(1), 87-131.

Pettibone, J. C. (2012). Testing the effect of time pressure on asymmetric dominance and compromise decoys in choice. *Judgment and Decision Making*, *7*(4), 513-523.

Piermont, E. (2017). Context dependent beliefs. *Journal of Mathematical Economics*, *71*, 63-73.

Piernas, C., & Popkin, B. M. (2011). Food portion patterns and trends among US children and the relationship to total eating occasion size, 1977–2006. *The Journal of Nutrition*, *141*(6), 1159-1164.

Pinger, P., Ruhmer-Krell, I., & Schumacher, H. (2016). The compromise effect in action: Lessons from a restaurant's menu. *Journal of Economic Behavior & Organization*, *128*, 14-34.

Plott, C. R. (1996). Rational individual behavior in markets and social choice processes: The discovered preference hypothesis. In C. Schmidt, E. Colombatto, & K. Arrow, (Eds.), *The Rational Foundations of Economic Behaviour* (pp. 225–250). London: McMillan.

Plott, C. R., & Zeiler, K. (2005). The willingness to pay-willingness to accept gap, the" endowment effect," subject misconceptions, and experimental procedures for eliciting valuations. *American Economic Review*, *95*(3), 530-545.

Pocheptsova, A., Amir, O., Dhar, R., & Baumeister, R. F. (2009). Deciding without resources: Resource depletion and choice in context. *Journal of Marketing Research*, *46*(3), 344-355.

Posner, M. I., & Cohen, Y. (1984). Components of visual orienting. In H. Bouma & D. G. Bouwhuis (Eds.), *Attention and performance* (vol. 10, pp. 531-556). Hillsdale, NJ: London.

Posner, M. I., Rafal, R. D., Choate, L. S., & Vaughan, J. (1985). Inhibition of return: Neural basis and function. *Cognitive Neuropsychology*, *2*(3), 211-228.

Pratkanis, A. R., & Farquhar, P. H. (1992). A brief history of research on phantom alternatives: Evidence for seven empirical generalizations about phantoms. *Basic and Applied Social Psychology*, *13*(1), 103-122.

Prelec, D., Wernerfelt, B., & Zettelmeyer, F. (1997). The role of inference in context effects: Inferring what you want from what is available. *Journal of Consumer Research*, *24*(1), 118-125.

Proto, E., & Sgroi, D. (2017). Biased beliefs and imperfect information. *Journal of Economic Behavior & Organization*, *136*, 186-202.

Radner, R., & Marschak, J. (1954). Note on some proposed decision criteria. In R. M. Thrall, C. H. Coombs, & R. L. Davies, (Eds.), *Decision Process* (pp. 61–68). New York: John Wiley & Sons, Inc.

Ray, P. (1973). Independence of irrelevant alternatives. *Econometrica*, *41*(5), 987-991.

Richter, M. K. (1966). Revealed preference theory. *Econometrica*, *34*(3), 635-645.

Riefer, P. S., Prior, R., Blair, N., Pavey, G., & Love, B. C. (2017). Coherency-maximizing exploration in the supermarket. *Nature Human Behaviour*, *1*(17), 1-4.

Rieskamp, J., Busemeyer, J. R., & Mellers, B. A. (2006). Extending the bounds of rationality: Evidence and theories of preferential choice. *Journal of Economic Literature*, *44*(3), 631-661.

Rippin, H. L., Hutchinson, J., Jewell, J., Breda, J. J., & Cade, J. E. (2019). Portion size of energy-dense foods among French and UK adults by BMI status. *Nutrients*, *11*(1), 12-32.

Riskey, D. R., Parducci, A., & Beauchamp, G. K. (1979). Effects of context in judgments of sweetness and pleasantness. *Perception & Psychophysics*, *26*(3), 171-176.

Robbins, L. (1932). *An Essay on the Nature and Significance of Economic Science*. London: Macmillan & Co..

Robinson, E., & Kersbergen, I. (2018). Portion size and later food intake: Evidence on the “normalizing” effect of reducing food portion sizes. *The American Journal of Clinical Nutrition*, *107*(4), 640-646.

Robinson, E., Henderson, J., Keenan, G. S., & Kersbergen, I. (2019). When a portion becomes a norm: Exposure to a smaller vs. larger portion of food affects later food intake. *Food Quality and Preference*, *75*, 113-117.

Robinson, E., Oldham, M., Cuckson, I., Brunstrom, J. M., Rogers, P. J., & Hardman, C. A. (2016). Visual exposure to large and small portion sizes and perceptions of portion size normality: Three experimental studies. *Appetite*, *98*, 28-34.

Robinson, E., te Raa, W., & Hardman, C. A. (2015). Portion size and intended consumption. Evidence for a pre-consumption portion size effect in males? *Appetite*, *91*, 83-89.

Rolls, B. J. (2014). What is the role of portion control in weight management? *International Journal of Obesity*, *38*(1), S1-S8.

Rolls, B. J., Morris, E. L., & Roe, L. S. (2002). Portion size of food affects energy intake in normal-weight and overweight men and women. *The American Journal of Clinical Nutrition*, *76*(6), 1207-1213.

Rolls, B. J., Roe, L. S., Kral, T. V., Meengs, J. S., & Wall, D. E. (2004). Increasing the portion size of a packaged snack increases energy intake in men and women. *Appetite*, *42*(1), 63-69.

Ronayne, D., & Brown, G. D. A. (2017). Multi-attribute decision by sampling: An account of the attraction, compromise and similarity effects. *Journal of Mathematical Psychology*, *81*, 11-27.

Rozin, P., & Schiller, D. (1980). The nature and acquisition of a preference for chili pepper by humans. *Motivation and Emotion*, *4*(1), 77-101.

Rulli, T., & Worsnip, A. (2016). IIA, rationality, and the individuation of options. *Philosophical Studies*, *173*(1), 205-221.

Samuelson, P. A. (1938). A note on the pure theory of consumer's behaviour. *Economica*, *5*(17), 61-71.

Samuelson, P. A. (1948). Consumption theory in terms of revealed preference. *Economica*, *15*(60), 243-253.

Sartori, N. (2006). Bias prevention of maximum likelihood estimates for scalar skew normal and skew t distributions. *Journal of Statistical Planning and Inference*, *136*(12), 4259-4275.

Satel, J., Wang, Z., Trappenberg, T. P., & Klein, R. M. (2011). Modeling inhibition of return as short-term depression of early sensory input to the superior colliculus. *Vision Research*, *51*(9), 987-996.

Sevilla, J., Lu, J., & Kahn, B. E. (2019). Variety seeking, satiation, and maximizing enjoyment over time. *Journal of Consumer Psychology*, *29*(1), 89-103.

Schulz, E., & Gershman, S. J. (2019). The algorithmic architecture of exploration in the human brain. *Current Opinion in Neurobiology*, *55*, 7-14.

Schulz, E., Bhui, R., Love, B. C., Brier, B., Todd, M. T., & Gershman, S. J. (2019). Structured, uncertainty-driven exploration in real-world consumer choice. *Proceedings of the National Academy of Sciences, 116*(28), 13903-13908.

Schulz, E., Wu, C. M., Ruggeri, A., & Meder, B. (2019). Searching for rewards like a child means less generalization and more directed exploration. *Psychological Science, 30*(11), 1561-1572.

Sen, A. K. (1973). Behaviour and concept of preference. *Economica*, *40*(159), 241-259.

Sen, A. K. (1977). Rational fools: A critique of the behavioral foundations of economic theory. *Philosophy & Public Affairs*, *6*(4), 317-344.

Sen, A. K. (1993). Internal consistency of choice. *Econometrica*, *61*(3), 495-521.

Settle, J. E., Dawes, C. T., & Fowler, J. H. (2009). The heritability of partisan attachment. *Political Research Quarterly*, *62*(3), 601-613.

Shadlen, M. N., & Shohamy, D. (2016). Decision making and sequential sampling from memory. *Neuron*, *90*(5), 927-939.

Shafir, S., Waite, T. A., & Smith, B. H. (2002). Context-dependent violations of rational choice in honeybees (Apis mellifera) and gray jays (Perisoreus canadensis). *Behavioral Ecology and Sociobiology*, *51*(2), 180-187.

Sharpe, K. M., Staelin, R., & Huber, J. (2008). Using extremeness aversion to fight obesity: Policy implications of context dependent demand. *Journal of Consumer Research*, *35*(3), 406-422.

Sharot, T., de Martino, B., & Dolan, R. J. (2009). How choice reveals and shapes expected hedonic outcome. *Journal of Neuroscience*, *29*(12), 3760-3765.

Sharot, T., Fleming, S. M., Yu, X., Koster, R., & Dolan, R. J. (2012). Is choice-induced preference change long lasting? *Psychological Science*, *23*(10), 1123-1129.

Sharot, T., Velasquez, C. M., & Dolan, R. J. (2010). Do decisions shape preference? Evidence from blind choice. *Psychological Science*, *21*(9), 1231-1235.

Sheng, S., Parker, A. M., & Nakamoto, K. (2005). Understanding the mechanism and determinants of compromise effects. *Psychology & Marketing*, *22*(7), 591-609.

Shogren, J. F., Margolis, M., Koo, C., & List, J. A. (2001). A random nth-price auction. *Journal of Economic Behavior & Organization*, *46*(4), 409-421.

Smith, G. E., & Nagle, T. T. (1995). Frames of reference and buyers' perception of price and value. *California Management Review*, *38*(1), 98-116.

Simonson, I. (1989). Choice based on reasons: The case of attraction and compromise effects. *Journal of Consumer Research*, *16*(2), 158-174.

Simonson, I. (1993). Get closer to your customers by understanding how they make choices. *California Management Review*, *35*(4), 68-84.

Simonson, I. (2008). Regarding inherent preferences. *Journal of Consumer Psychology*, *18*(3), 191-196.

Simonson, I. (2014). Vices and virtues of misguided replications: The case of asymmetric dominance. *Journal of Marketing Research*, *51*(4), 514-519.

Simonson, I., Bettman, J. R., Kramer, T., & Payne, J. W. (2013). Comparison selection: An approach to the study of consumer judgment and choice. *Journal of Consumer Psychology*, *23*(1), 137-149.

Simonson, I., & Nowlis, S. M. (2000). The role of explanations and need for uniqueness in consumer decision making: Unconventional choices based on reasons. *Journal of Consumer Research*, *27*(1), 49-68.

Simonson, I., & Sela, A. (2011). On the heritability of consumer decision making: An exploratory approach for studying genetic effects on judgment and choice. *Journal of Consumer Research*, *37*(6), 951-966.

Simonson, I., & Tversky, A. (1992). Choice in context: Tradeoff contrast and extremeness aversion. *Journal of Marketing Research*, *29*(3), 281-295.

Slovic, P. (1995). The construction of preference. *American Psychologist*, *50*(5), 364-371.

Smirnov, I., & Thurner, S. (2017). Formation of homophily in academic performance: Students change their friends rather than performance. *PloS One*, *12*(8), e0183473.

Smith, E. R. (2008). Embodied constraints as elements in attitude construction. *Journal of Consumer Psychology*, *18*(3), 187-190.

Sokol-Hessner, P., & Rutledge, R. B. (2019). The psychological and neural basis of loss aversion. *Current Directions in Psychological Science*, *28*(1), 20-27.

Spektor, M. S., Gluth, S., Fontanesi, L., & Rieskamp, J. (2019). How similarity between choice options affects decisions from experience: The accentuation-of-differences model. *Psychological Review*, *126*(1), 52-88.

Spence, M. (2002). Signaling in retrospect and the informational structure of markets. *American Economic Review*, *92*(3), 434-459.

Steenhuis, I. H., & Vermeer, W. M. (2009). Portion size: Review and framework for interventions. *International Journal of Behavioral Nutrition and Physical Activity*, *6*(58), 1-10.

Steenhuis, I. H., Leeuwis, F. H., & Vermeer, W. M. (2010). Small, medium, large or supersize: Trends in food portion sizes in The Netherlands. *Public Health Nutrition*, *13*(6), 852-857.

Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology*, *53*(1), 1-26.

Stojić, H., Schulz, E., P Analytis, P., & Speekenbrink, M. (2020). It's new, but is it good? How generalization and uncertainty guide the exploration of novel options. *Journal of Experimental Psychology: General*, *149*(10), 1878-1907.

Taylor, J. E. T., Chan, D., Bennett, P. J., & Pratt, J. (2015). Attentional cartography: mapping the distribution of attention across time and space. *Attention, Perception, & Psychophysics*, *77*(7), 2240-2246.

Taylor, T. L., & Klein, R. M. (1998). On the causes and effects of inhibition of return. *Psychonomic Bulletin & Review*, *5*(4), 625-643.

Thaler, R. (1992). *The winner's curse: Paradoxes and anomalies of economic life*. New York, NY: Cambridge University Press.

Theeuwes, J. (2010). Top–down and bottom–up control of visual selection. *Acta Psychologica*, *135*(2), 77-99.

Tipper, S. P., Weaver, B., Jerreat, L. M., & Burak, A. L. (1994). Object-based and environment-based inhibition of return of visual attention. *Journal of Experimental Psychology: Human Perception and Performance*, *20*(3), 478-499.

Tufano, F. (2010). Are 'true'preferences revealed in repeated markets? An experimental demonstration of context-dependent valuations. *Experimental Economics*, *13*(1), 1-13.

Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review*, *79*(4), 281-299.

Tversky, A., & Simonson, I. (1993). Context-dependent preferences. *Management Science*, *39*(10), 1179-1189.

Uysal, S. D., & Pohlmeier, W. (2011). Unemployment duration and personality. *Journal of Economic Psychology, 32*(6), 980-992.

van der Bend, D., Bucher, T., Schumacher, T. L., Collins, K., de Vlieger, N., Rollo, M., Burrows, T. L., Watson, J. F., & Collins, C. E. (2017). Trends in food and beverage portion sizes in Australian children; a time-series analysis comparing 2007 and 2011–2012 national data. *Children*, 4(8), 69-77.

Vazire, S., & Mehl, M. R. (2008). Knowing me, knowing you: the accuracy and unique predictive validity of self-ratings and other-ratings of daily behavior. *Journal of Personality and Social Psychology, 95*(5), 1202-1216.

Vermote, M., Versele, V., Stok, M., Mullie, P., D'Hondt, E., Deforche, B., Clarys, P., & Deliens, T. (2018). The effect of a portion size intervention on French fries consumption, plate waste, satiety and compensatory caloric intake: An on-campus restaurant experiment. *Nutrition Journal*, *17*(43), 1-12.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, *16*(1), 8-37.

Vlaev, I., Chater, N., Stewart, N., & Brown, G. D. A. (2011). Does the brain calculate value? *Trends in Cognitive Sciences*, *15*(11), 546-554.

Voigt, K., Murawski, C., & Bode, S. (2017). Endogenous formation of preferences: Choices systematically change willingness-to-pay for goods. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *43*(12), 1872-1882.

Wansink, B., & van Ittersum, K. (2007). Portion size me: Downsizing our consumption norms. *Journal of the American Dietetic Association, 107*(7), 1103-1106.

Wansink, B., Painter, J. E., & North, J. (2005). Bottomless bowls: Why visual cues of portion size may influence intake. *Obesity Research*, *13*(1), 93-100.

Warren, C., McGraw, A. P., & Van Boven, L. (2011). Values and preferences: Defining preference construction. *Wiley Interdisciplinary Reviews: Cognitive Science*, *2*(2), 193-205.

Weber, E. U., & Johnson, E. J. (2009). Mindful judgment and decision making. *Annual Review of Psychology*, *60*, 53-85.

Weiss-Cohen, L., Konstantinidis, E., Speekenbrink, M., & Harvey, N. (2016). Incorporating conflicting descriptions into decisions from experience. *Organizational Behavior and Human Decision Processes*, *135*, 55-69.

Wernerfelt, B. (1995). A rational reconstruction of the compromise effect: Using market data to infer utilities. *Journal of Consumer Research*, *21*(4), 627-633.

Wilson, T. D. (2002). Strangers to ourselves: Discovering the adaptive unconscious. Cambridge, MA, US: Belknap Press/Harvard University Press.

Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., & Cohen, J. D. (2014). Humans use directed and random exploration to solve the explore–exploit dilemma. *Journal of Experimental Psychology: General*, *143*(6), 2074-2081.

Wrieden, W., Gregor, A., & Barton, K. (2008). Have food portion sizes increased in the UK over the last 20 years?. *The Proceedings of the Nutrition Society*, *67*(OCE).

Wu, C., & Cosguner, K. (2020). Profiting from the Decoy Effect: A case study of an online diamond retailer. *Marketing Science*, *39*(5), 974-995.

Wu, C. M., Schulz, E., Speekenbrink, M., Nelson, J. D., & Meder, B. (2018). Generalization guides human exploration in vast decision spaces. *Nature Human Behaviour*, *2*(12), 915-924.

Wu, L., Gong, X., Chen, X., & Hu, W. (2020). Compromise effect on food consumer choice behaviors in China. *Frontiers in Psychology*, *11*(1352), 1-10.

Yap, B. W., & Sim, C. H. (2011). Comparisons of various types of normality tests. *Journal of Statistical Computation and Simulation*, *81*(12), 2141-2155.

Young, L. R., & Nestle, M. (2002). The contribution of expanding portion sizes to the US obesity epidemic. *American Journal of Public Health*, *92*(2), 246-249.

Young, L. R., & Nestle, M. (2007). Portion sizes and obesity: Responses of fast-food companies. *Journal of Public Health Policy*, *28*(2), 238-248.

Zell, E., & Krizan, Z. (2014). Do people have insight into their abilities? A metasynthesis. *Perspectives on Psychological Science*, *9*(2), 111-125.

Zeman, A. (2001). Consciousness. *Brain*, *124*, 1263-1289.

Zeman, A. (2006). What do we mean by “conscious” and “aware”?. *Neuropsychological Rehabilitation*, *16*(4), 356-376.

Zlatevska, N., Dubelaar, C., & Holden, S. S. (2014). Sizing up the effect of portion size on consumption: A meta-analytic review. *Journal of Marketing*, *78*(3), 140-154.