Issues facing the modern consumer: topics in industrial organisation and decision-making

by

David Ronayne

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

Department of Economics

September 2015

THE UNIVERSITY OF
WARWICK

Contents

Acknowledgements 1
Declaration 2
Abstract 3

Chapter 1. Price Comparison Websites
- Introduction 4
- Literature 7
- Model 9
- Comparative Statics 14
- Competing Aggregators 17
- Publicly-Observable Fees 21
- Price Discrimination 22
- The Extensive Search Margin 24
- Conclusion 27
- Appendix 28
- References 49

Chapter 2. Multi-Attribute Decision by Sampling: An Account of the Attraction, Compromise and Similarity Effects
- Introduction 55
- Context Effects 57
- Literature 58
- Multi-Attribute Decision by Sampling 60
- Experiments 66
- Discussion 79
- Appendix A - Methodological Details 82
- Appendix B - Sufficient Conditions for Context Effects 84
- References 92

Chapter 3. E-Cigarettes: The Extent and Impact of Complementary Dual-Use
- Introduction 100
- Model 102
- The Survey 105
- Cost Benefit Analysis 113
- Conclusion 120
- References 121
- Supplementary Appendix 127

List of Figures

Chapter 1. Price Comparison Websites

Figure 1: Equilibrium price distributions 15

Chapter 2. Multi-Attribute Decision by Sampling: An Account of the Attraction, Compromise and Similarity Effects

Figure 1: Context effects 58

Figure 2: Multi-Attribute Decision by Sampling 62

Figure 3: Attraction effect 64

Figure 4: Compromise effect 65

Figure 5: Similarity effect 66

Figure 6: Screen-shot of manipulation 68

Figure 7: Manipulating choice: method 69

Figure 8: Manipulating context effects: method 71

Figure 9: Manipulating context effects: evidence 71

Figure 10: Screen-capture of distribution elicitation 74

Figure 11: Plotting data 76

Figure A1: Hotel data 82

Figure A2: Manipulating context effects: evidence (all) 83

Chapter 3. E-Cigarettes: The Extent and Impact of Complementary Dual-Use

Figure 1: Screen-shot of the Substitute or Complement Question 107

Figure 2: Empirical CDF of Responses to the Question of Figure 1 108

Figure 3: Empirical CDF Decomposed by Quit Status 110

List of Tables

Chapter 2: Multi-Attribute Decision by Sampling: An Account of the Attraction, Compromise and Similarity Effects

Table 1: Manipulating choice 70

Table 2: Predicting individuals choices 78

Table A1: Choice sets 83

Chapter 3: E-Cigarettes: The Extent and Impact of Complementary Dual-Use

Table 1: Respondent Demographics 106

Table 2: The Perception Gap 109

Table 3: Characteristics of Complementary and Substitutive Dual-Users 112

Table 4: Cost-Benefit Analysis Scenarios 119

Acknowledgements

I thank my supervisors Motty Perry and Daniel Sgroi for their assistance and guidance throughout the PhD. I also thank Dan Bernhardt for his help and time over the final two years of the program. Chapter-specific acknowledgements are as follows:

Price Comparison Websites

I thank Dan Bernhardt, Rani Spiegler, Motty Perry, Daniel Sgroi, Kobi Glazer, Andrew Oswald, Alessandro Iaria and Giulio Trigilia for their helpful discussion. I also thank participants of the RES and EARIE 2015 conferences alongside those at various other seminars.

Multi-Attribute Decision by Sampling: An Account of the Attraction, Compromise and Similarity Effects

I thank Gordon Brown (co-author), Neil Stewart, Takao Noguchi, Konstantinos Tsetsos, Daniel Sgroi, Andrew Oswald, Sudeep Bhatia, Timothy Mullett, Lukasz Walasek and Mario Sanclemente, participants of 2014 Bristol Decision-Making Conference and the Decision Research at Warwick Forum for helpful comments and discussion. I thank Andrew Taylor for his programming and software services. The work was supported by the Economic and Social Research Council (grant number ES/K002201/1) and the Leverhulme Trust (grant number RP2012-V-022).

E-Cigarettes: The Extent and Impact of Complementary Dual-Use

I thank Daniel Sgroi (co-author), Chris Doyle (co-author), Monica Giovanniello, Abhinay Muthoo and members of the Warwick Policy Lab for their discussion and comments. Funding was provided by the Warwick Policy Lab. Ethical approval for the survey was obtained from the Humanities and Social Sciences Research Ethics Committee at the University of Warwick on 21 November 2014 under reference 46/14-15.

Declarations

This thesis has not been submitted for a degree at another university. Any errors in this thesis are my responsibility. The chapters are products of varying levels of collaboration, as follows: *Price Comparison Websites*, solely my own work; *Multi-Attribute Decision by Sampling: An Account of the Attraction, Compromise and Similarity Effects*, co-authored with Gordon Brown (Department of Psychology, University of Warwick); *E-Cigarettes: The Extent and Impact of Complementary Dual-Use*, co-authored with Daniel Sgroi (my supervisor) and Chris Doyle (Department of Economics, University of Warwick). All data analysis, figures and tables in this thesis were composed my me.

Abstract

This thesis deals with specific topics within industrial organisation and consumer choice. The first chapter *"Price Comparison Websites"* deals with a new and growing multi-billion dollar industry that has emerged and flourished over the last two decades. Price comparison websites (PCWs) or 'web aggregators' are poised to benefit consumers by increasing competitive pricing pressure on firms by acquainting shoppers with more prices. However, these sites also charge firms for sales which feeds back to raise prices. I find that the introduction of a PCW *increases* prices for all consumers. I then consider market outcomes with competing PCWs, fee-transparency, price discrimination and costly search. The findings suggest the industry may in fact *not* benefit any consumers, regardless of whether they actually use the sites, and that regulatory solutions are not as simple as they may seem.

The second chapter *"Multi-Attribute Decision by Sampling: An Account of the Attraction, Compromise and Similarity Effects"* (MADS) proposes and tests a theory that offers an explanation of the three major puzzles, known as 'context effects', observed in multi-attribute choice: the attraction, compromise and similarity effects. Existing models within the judgement and decision-making (JDM) literature all propose complex choice architectures that cannot explain all three effects for the same reason and in most cases, the models are only computationally solvable. Our approach is based on a single mechanism, able to explain all three effects simultaneously and permits analytic expressions for choice probabilities. MADS combines and extends three independently-developed ideas from psychology: (i) that individuals compare choice alternatives with comparison items; (ii) that comparisons are based on simple dominance relations; and (iii) that comparison items are systematically influenced by the choice set faced. In a novel experimental design involving 1,200 online participants and real-world data on hotel stays, we find support for our theory's account. In fact, our treatment was strong enough to make the classic context effects appear and disappear. It is concluded that a single-mechanism model, Multi-Attribute Decision by Sampling, can account for the three classic context effects.

In the third chapter *"E-Cigarettes: The Extent and Impact of Complementary Dual-Use"*, we offer one of the first studies of the electronic cigarette market from an economic perspective. We apply the classic economic notion of complements and substitutes to electronic and conventional cigarettes. The work examines the validity of the common assumption that smokers adopt e-cigarettes to reduce (i.e., substitute) their consumption of traditional cigarettes. We ran an online survey to assess the motivations of users of both products, referred to as dual-users. The results show 37% of dual-users regard e-cigarettes primarily as a *complement* to conventional cigarettes. Using our survey together results alongside publicly-available US data, we calibrate a cost-benefit analysis and show that complementary use could wipe out up to 57% of the financial benefits e-cigarettes offer to society.

Price Comparison Websites

1. Introduction

Over the past two decades a new industry of price comparison websites (PCWs) or 'web aggregators' has emerged. The industry has enabled consumers to check prices of many firms selling a particular service or product simultaneously in one place. The sites are popular in many countries, and in many markets including utilities, financial services, hotels, flights and durable goods.[1] These sites command billions of dollars of revenue annually.[2] In the UK, PCWs for utilities and financial services have been particularly successful. There are roughly 48 such PCWs in the UK, where over 70% of internet users have used such a site. The largest four aggregators generated approximately £800m ($1.2bn) in revenue during 2013, with average annual profit of the group increasing by 14% that year.[3]

The internet has altered search costs, allowing consumers to compare prices across firms in a matter of clicks, intensifying competitive pricing pressure between firms. While a consumer may not know all the firms in a market, a PCW can expose the full list of market offerings, maximizing inter-firm pressure. However, underlying this increased competition are the fees paid by firms who sell their products through the websites. In the UK these are understood to be between £44-60 ($69-95) for a customer switching gas and electricity provider.[4] These fees, in turn, act as a marginal cost faced by producers, affecting their pricing decisions.[5]. The industry gleans substantial profits

[1]Examples for utilities and services include Moneysupermarket.com, Google Compare and Gocompare.com; for flights Skyscanner.net and Flights.com; for hotels Hotels.com and Booking.com; and for durable goods Amazon Marketplace, Pricerunner.com and Pricegrabber.com.

[2]Regarding travel services, Priceline Group (who own Booking.com and Priceline.com) and Expedia Inc. (who own Expedia.com and Hotels.com) made approximately $6bn in total agency revenues in 2014. Regarding durable goods, Amazon Marketplace sold 2 billion items from third-party sellers. See their 2014 Annual Reports for details.

[3]Number of PCWs taken from Consumer Focus (2013) report into PCWs in the UK. The 'big four' refers to Money Supermarket, Compare the Market, Go Compare and Confused.com. Number of PCWs Usage data from the 2013 Mintel Report on Web Aggregators. Financial information taken from the companies' own annual reports where available, otherwise inferred from parent group reports or newspaper articles, the details are in the appendix.

[4]For example, see BBC (2015). Fees are significant in other sectors too e.g., in the hotel-reservation sector they are reported to be 15-25% of the purchase price (see Daily Mail 2015).

[5]In practice, PCWs in some markets charge per-click, some per-sale. I abstract from this difference

from these fees. As such, it is not clear whether the central premise that PCWs lower prices is valid. This fundamental tension is encapsulated in a quote from the BBC (2014):

> "There's another cost in the bill. It's hidden, it's kept confidential, and yet it's for a part of the industry that appears to be on the consumers' side. This is the cut of the bill taken by price comparison websites, in return for referring customers. The recommendation to switch creates churn in the market, and it is seen by supplier companies as worth paying high fees to the websites. Whether or not customers choose to use the sites, the cost to the supplier is embedded within bills for all customers." [6]

This article examines this "churn" and addresses the fundamental question of when consumers are better off with a PCW in the marketplace. In homogeneous good markets, I characterize when all consumers are made worse off following the introduction of PCWs.

My model builds on the elegant framework of Baye and Morgan (2001), who investigate the strategic incentives of a PCW or 'information gatekeeper'. Without the PCW in their model, each consumer is served by a *single* 'local' firm which sells at the monopoly price (it is too costly for consumers to travel to another store). This leads to the result that consumers benefit from the introduction of a PCW, because firms must compete for the business of consumers who enjoy free access to the site. In the modern online marketplace however, firms also have their own websites.[7] In the absence of an aggregator, consumers do not need to physically travel to purchase the good; they can visit another firm's website just as easily as they could an aggregator's.

My model features two types of consumers: shoppers, who use PCWs in equilibrium, and inactive consumers, who buy directly from a particular firm. Although shop-

---

and model the latter, assuming that the two are correlated.

[6] The concern was also expressed by US senator Amy Klobuchar regarding mergers of hotel-reservation sites: "The whole idea of cheaper hotels is very good, but if it all starts to come under one company, you can easily foresee the situation where they can charge higher commissions that are then passed on to consumers." (New York Times, 2015)'

[7] Often, a PCW simply re-directs users to the selected firm's own website to complete the purchase.

pers are always better off than inactive consumers in equilibrium, my primary result is that *both* types are always made worse off by the introduction of a single PCW. I then provide conditions under which consumers are harmed by the introduction of multiple competing PCWs. This is the first article in this setting to show such results, reversing those in the existing literature, which I show can be seen as a special limiting case.[8] My model supposes each shopper is informed of the prices from $q > 1$ of the n firms' websites in the absence of a PCW, rather than $q = 1$ as in Baye and Morgan. Without a PCW, shoppers see q prices; after the introduction of the PCW, they will see all of them in equilibrium. I show that adding competition ($q > 1$) to the setting without an aggregator reverses their finding: expected equilibrium prices are raised by the introduction of a single PCW, making all consumers worse off.

What happens is that the equilibrium fee a single PCW charges for a sale through its site is so high that it more than negates any benefits from the increased firm competition. One may conjecture that allowing multiple competing aggregators will undo this result, akin to textbook Bertrand competition. I extend the model to allow for multiple PCWs, and for shoppers to check any number of them. My characterization shows that both the number of PCWs and the number of PCWs that shoppers check (or are informed of), matter. Concretely, when shoppers only check one of many PCWs, they all effectively remain monopolists and so consumer welfare does improve with any number of PCWs. At the other extreme where shoppers check all PCWs, Bertrand-style reasoning at the aggregator level results as a special case: PCWs undercut each other's fees to reach a unique zero-profit equilibrium and shoppers benefit from their existence. When shoppers some, but not all PCWs, consumer welfare tends to rise with the number of PCWs checked, but can fall with the number of PCWs. In particular, in this realistic case, there is a critical number of PCWs beyond which all consumers can be worse off than without any aggregators at all.

I then investigate the impact of different policies and other market features. Generally, PCW fees are either not disclosed or are detailed on a subpage of the websites. I

[8]This is in the absence of any persuasion or direction of consumers to more expensive products by firms (e.g., Armstrong and Zhou, 2011), or biased intermediaries (e.g., de Cornière and Taylor, 2014).

extend the analysis to settings where PCWs publicly announce fees so that consumers are made aware of them. This creates the possibility that firms and shoppers could coordinate on the cheapest PCW, which can result in equilibria that benefit shoppers relative to a world without PCWs.

My primary focus is on settings where a firm sets the same price for a product on every website that it is sold. This assumption describes many markets where PCWs operate including gas, electricity, mortgages and durable goods. However, I extend the analysis to consider markets such as the hotel reservation sector, where a firm's price can differ across sites. In this case, I show that both types of consumers can be worse off when shoppers only check one PCW in equilibrium. However, unlike markets without price discrimination, when shoppers check at least two PCWs, competitive pressure between aggregators can work à la Bertrand. I also consider markets where consumers may face some non-negligible search cost in order to retrieve prices e.g., home insurance quotes. Here, the number of shopping consumers is endogenously determined by the distribution of search costs and the expected savings consumers can make by using it. I show that despite this activity at the 'extensive search margin', an aggregator can still make all consumers worse off, including those who decide to start shopping.

Section 2 reviews the literature; Section 3 presents the model; Section 4 conducts comparative statics with a monopolist PCW; Section 5 models competing aggregators; Sections 6-8 consider settings with publicly observable fees, price discrimination and search costs; Section 9 concludes.

2. Literature

This article contributes to the literature on 'clearing-house' models, see for example (Salop and Stiglitz, 1977; Varian, 1980; Rosenthal, 1980; Baye and Morgan, 2001, 2009; Baye et al., 2004; Chioveanu, 2008; Arnold et al., 2011; Arnold and Zhang, 2014). These models rationalize price-dispersion in homogeneous goods markets.[9] In-

[9]The search cost literature also provides explanations of price dispersion (e.g., Burdett and Judd, 1983; Ellison and Ellison, 2009; Ellison and Wolitzky, 2012; Stahl, 1989; 1996; Stigler, 1961). Motivated

deed, price dispersion has persisted despite the advances of technology such the internet and comparison sites. Early studies documented marked dispersion in the online markets for goods (e.g., Brynjolfsson and Smith, 2000; Baye et al., 2004). A recent study by Gorodnichenko et al. (2015) finds substantial cross-seller variation in prices, and voices support for clearing-house models that categorize consumers into loyal and shopping consumers. The equilibria in my model feature price dispersion regardless of whether there is an aggregator (or 'clearing-house'). Without a PCW, this is because there is *some* consumer search. Producing price dispersion with a PCW that employs the pricing mechanism seen in practice, is a challenge. The shift in the aggregator industry away from charging one-off fixed fees, toward pay-per-sale fees to firms is rationalized as profit-maximizing PCW behavior by Baye et al. (2011). However, without the introduction of some other exogenous fixed cost to a firm of listing on the PCW (e.g., transaction costs), price-dispersion vanishes in equilibrium. I do not deny the existence of such additional costs, but emphasize that dispersion arises in my framework without appeal to fixed listing costs.

The larger relevant literature is that of two-sided markets, pioneered by Rochet and Tirole (2003) (see also Caillaud and Jullien, 2003; Ellison and Fudenberg, 2003; Armstrong, 2006; Reisinger, 2014). These articles model platforms where buyers and sellers meet to trade, focusing on platform pricing and the effect of network externalities with differentiated products and platforms. These models do not explicitly model seller-side competition, which is central to my model. More recent contributions do (e.g., Belleflamme and Peitz, 2010; Hagiu, 2009) but they model the platform as the only available technology, which is not appropriate for the questions I address. Edelman and Wright (2015) allow sellers and buyers to conduct business off platform, but they study environments with differentiated product markets, where the intermediary directly offers buyer-side benefits such as rebates. In contrast, I model a homogeneous good, isolating price as the determinant of consumer welfare, where the only benefit a 'platform' brings is informational: it lists available prices. The potential benefit of a PCW to consumers is

by the rise of the internet, clearing-house models abstract from a direct modeling of consumer search costs. The frameworks are to some extent isomorphic (see Baye et al., 2006).

that it can lower prices through the interaction of strategic, competing firms, and hence the size of the benefit is determined endogenously via the equilibrium actions of firms and consumers.

3. Model

A World Without a PCW

There are n firms and a unit-mass of consumers. Firms produce a homogeneous product at zero cost without capacity constraints. Consumers wish to buy one unit and have a common willingness to pay of $v > 0$. Each consumer is endowed with a 'default', 'current' or 'preferred' firm from which they are informed of the price. This assumption has many natural interpretations. In the market for services and utilities (e.g., gas and electricity tariffs, mortgages, credit cards, broadband, cellphone contracts, car, home and travel insurance etc.) consumers can be thought of as having a current provider for the service for which they know the price they pay and the renewal price should they remain with the same provider.[10] In the market for flights, hotels or durable goods, consumers can be thought to have a carrier, hotel (or hotel chain) or producer which they prefer, perhaps bookmarked in their browser or for which they receive marketing emails that directly inform them of the price or prompt the user to find out via hyperlinks, in a matter of seconds.

A proportion $\alpha \in (0, 1)$ of consumers are 'shoppers'. Casual empiricism suggests that many people enjoy browsing or looking for a bargain. Shoppers check (are aware of) at least one rival firm's website and therefore know $q > 1$ of the n prices. A shopper sticks with his or her default firm if its price is not beaten by another price. If a rival firm's price is cheaper, the shopper switches. The remaining proportion of $1 - \alpha$ consumers are 'auto-renewing', 'loyal', 'offline' or 'inactive' consumers who do not shop around. In markets where firms are service providers, such consumers simply allow their contract with their existing provider to continue with their default firm. Much of the furore surrounding PCWs has been directed at those operating in the

[10] Providers usually provide renewal price information directly to consumers.

services and utilities sectors. In the presentation that follows I adopt terminology suited to that sector, of 'current' firms or providers with consumers referred to as 'shoppers' and 'auto-renewers'.

I assume that each firm has an equal share of current consumers of each type ($\frac{\alpha}{n}$ shoppers and $\frac{1-\alpha}{n}$ auto-renewers). Firms set prices and shoppers simultaneously decide which rival-firm websites to visit. I focus on equilibria in which firms adopt identical pricing strategies and shoppers employ symmetric shopping strategies. Going forward I refer to such symmetric equilibria simply as equilibria.

This setup generalizes Varian (1980), nesting his equilibrium.[11] He motivates the two types as being completely 'informed' about all n prices, or 'uninformed'. The informed buy the cheapest on the market, while the uninformed buy from their default firm. In my model, shoppers see q prices where $1 < q \leq n$ and I derive the unique equilibrium. When $q = n$, Varian's equilibrium corresponds to that derived in Proposition 1 below.

Shoppers can be characterized by $\binom{n}{q}$ groups. Each group is a list of the firms checked by that consumer. For example, consider $q = 2$, $n = 4$ with firms indexed 1,2,3,4. Then there are 6 possible comparisons shoppers could make: $\{$12, 13,14, 23, 24, 34$\}$, where the two digits refer to which firms' prices are checked. Each firm is involved in $\binom{n-1}{q-1} = 3$ price comparisons. Shoppers employ symmetric shopping strategies, and hence are evenly distributed across these groups: $\frac{1}{6}$ compare the prices of firms 1 and 2, $\frac{1}{6}$ compare firms 1 and 3, and so on. Equivalently, one could interpret a consumer as randomizing uniformly over which rival-firm websites she or he checks.[12]

Now consider the best-response of firms. Without loss of generality, let $F : p \rightarrow [0, 1]$ denote the cumulative distribution function of prices charged by firms in equilibrium. Proposition 1 describes the equilibrium without a PCW.

Proposition 1. *In the unique equilibrium firms mix according to the distribution*

[11] The only differences being that I assume zero costs and do not employ his zero-profit condition.

[12] Note that although I focus on a symmetric distribution across these pairs, this shopper behavior is not an assumption per se; it is a best response in equilibrium to symmetric firm behavior because consumers are then indifferent to which firm they check.

$$F(p) = 1 - \left[\frac{(v-p)(1-\alpha)}{qp\alpha}\right]^{\frac{1}{q-1}} \textit{ over the support } p \in \left[\frac{v(1-\alpha)}{1+\alpha(q-1)}, v\right].$$

Price dispersion is a central feature in clearing-house models and showcases how they rationalize price dispersion in homogeneous goods markets. The limited-search assumption does not alter the intuition for why. Bertrand-style reasoning of undercutting down to the marginal cost (here normalized to 0) does not play out. Because of the existence of auto-renewers, firms can guarantee themselves profit of at least $\frac{v(1-\alpha)}{n} > 0$. Therefore, any point mass existing in equilibrium strategies must be for some $\dot{p} > 0$. Any such mass would always be undercut by firms to gain a discrete number of shoppers for an arbitrary ϵ-loss in price.

Models building on Varian's assume that in a world without a clearing-house, consumers cannot check other prices, so the pure monopolistic-price equilibrium of $p = v$ results. As in Stahl (1989), if there were no shoppers in my model, I would obtain such a Diamond (1971) equilibrium with each firm charging v; and if there were only shoppers, then the Bertrand outcome of $p = 0$ would result. Unlike Stahl, shoppers do not necessarily know the prices of all firms, but they know at least two. Proposition 1 shows that with some (but non-zero) search, price dispersion still emerges in equilibrium.

It is also instructive to note that equilibrium pricing does not vary with the number of firms, n, as long as $q < n$. Each shopper compares q prices, regardless of the total number of firms. When making pricing decisions, a firm is not concerned about the number of other firms per se, but rather about the other prices shoppers know. As all firms price symmetrically and independently in equilibrium, it is as if each firm only faces $q-1$ rivals. In other words, what matters is the number of comparisons shoppers make, not the number of firms in the market.

A World With a PCW

Suppose an entrepreneur creates a price comparison website. I add a preliminary stage to the game at which the PCW sets a 'click-through fee' $c \in \mathbb{R}_+$ that a firm must pay to the aggregator per sale made via the site. Consumers do not learn c, although they will have correct expectations in equilibrium.[13] Each firm sees the fee c and must choose a price, and whether or not to post it on the PCW.

Models with a clearing-house that does not charge a fixed fee typically have many equilibria, see for instance Baye et al. (1992) for their characterization of the full set of equilibria of Varian (1980). Throughout the paper, I focus on symmetric equilibria in which shoppers only check PCWs.[14] With a monopoly PCW, this results in a unique equilibrium that one can explicitly solve for. In the equilibrium, firms list on the PCW with probability one and prices are dispersed, so shoppers have a strict incentive to check the PCW.[15]

To find the equilibrium, I use the following lemma which takes c as fixed and characterizes the ensuing mutual best-responses of firms. Define $G(p;c)$ to be the cumulative distribution function of prices charged by firms for a given click through fee c:

$$G(p;c) = 1 - \left[\frac{(v-p)(1-\alpha)}{\alpha(np-(n-1)c)}\right]^{\frac{1}{n-1}}$$

Lemma 1. *The mutual best-responses of firms as a function of c:*

1. *If $c \in [0, v(1-\alpha))$, firm responses have no point masses, and are described by $G(p;c)$.*

2. *If $c = v(1-\alpha)$, there are two classes of response, one with no point masses described by $G(p,c)$, and those that are degenerate.*

3. *If $c \in (v(1-\alpha), v]$, firm responses are degenerate.*

Where pricing is described by $G(p;c)$ each firm always lists its price on the PCW.

When c is in the lower interval described by Lemma 1, firms undercut each other until the point at which they would be better off charging the maximum price v and only

[13]As the BBC quote in the Introduction notes, the exact fee is "kept confidential", not publicly announced.

[14]Recall this is in addition to knowing their current or preferred firm's price, which upon sale, continues to count as a direct purchase e.g., the auto-renewal letter, the bookmarked airline or the hotel's marketing email.

[15]Note more generally that there exist other equilibria e.g., trivial equilibria in which no firms or shoppers attend the PCW. The set-up is then identical to that of the previous section and the equilibrium is given by Proposition 1 with the vacuous addition that the PCW can charge any c. There also exist asymmetric equilibria where not all firms list on the PCW (see Footnote 16).

selling to their auto-renewers. However, if c exceeds $v(1-\alpha)$, undercutting does not reach this point, so no firm would jump to v. Rather, it reaches a point where all firms charge the same price, selling to all their consumers directly, where no firm undercuts further or jumps to v. The reason no further undercutting occurs is that doing so would mean charging a price $p < c$ that would win all rival firms' shoppers, but would do so at a loss. The reason no firm jumps to v is that they are making more in this pure-pricing equilibrium where they sell to all their consumers directly. If $c = v(1-\alpha)$ then both mixed and pure equilibria obtain as threshold cases.

To derive the equilibrium, consider now incentives of the PCW. The PCW will make zero profit if firms all charge the same price, as then no shoppers switch, but positive profit in any mixing equilibrium if $c > 0$. The PCW therefore has a strong incentive to induce price dispersion because shoppers switch when they can obtain a strictly lower price with a new firm. A feature of distributions with no point masses is that the probability of a tie in price is zero. As a result, shoppers at $n-1$ of the n firms will switch. Given that firms mix in this way, the PCW will raise c as high as possible before reversion to a pure equilibrium, which here happens for $c > v\,(1-\alpha)$. Proposition 2 characterizes the equilibrium.[16]

Proposition 2. *In the unique equilibrium where shoppers check the PCW, the PCW sets a click-through fee of $c = v\,(1-\alpha)$, firms list on the PCW and mix over prices according to $G(p; v(1-\alpha))$ over the support $p \in [v(1-\alpha), v]$.*

That only pure equilibria exist for $c > v(1-\alpha)$ is what limits the PCW fee and allows there to be price dispersion at equilibrium without fixed costs for firms. It is the outside option of current shoppers for firms that limits the fee that the PCW charges. Shoppers know their current provider's price regardless of where they check prices and stay with their current firm if there are no lower prices to be found on the PCW. Hence,

[16] Note that in the symmetric equilibrium, all firms list on the PCW. In practice however, some firms do not always list on PCWs. Note that here there exist asymmetric equilibria where $m \geq 2$ firms list on the PCW, mixing over prices $[v(1-\alpha), v]$ in a similar way to the CDF of Proposition 2, with the other $n-m$ firms not listing, charging v and selling only to their auto-renewers with all shoppers uniformly spread over the mixing firms. The price-rising result of this article also holds in any of these asymmetric equilibria.

a firm can always guarantee itself its own shoppers and avoid the aggregator's fee by charging a price low enough to undercut the market and not list on the PCW. It is this threat that discourages the PCW from charging fees that are even higher in equilibrium. If firms did not have this outside option, the PCW could raise its fee to $c = v$, firms would charge $p = v$, price dispersion would be lost and the PCW would be able to extract all the surplus.

4. Comparative Statics

Comparing the equilibria of Propositions 1 and 2 reveals how the PCW affects consumer welfare.

Proposition 3. *Both types of consumer are worse off with the PCW than without.*

The proof of Proposition 3 shows that the expected shopper-price under F is less than the lower bound of the support of G (see Figure 1). It immediately follows that shoppers expect to pay more under G, as the expected lowest price is higher. That G first-order stochastic dominates F shows that auto-renewers expect to pay more under G, as the expected price is higher.

With the introduction of a PCW, the two effects stated in Corollaries 1 and 2 give rise to Proposition 3:

Corollary 1. *Within the mixed-price equilibrium firm responses of Lemma 1, as* $c \in [0, v(1-\alpha)]$ *increases, the expected price paid by both types of consumer increases.*

Corollary 2. *As the number of firms increases, the expected price paid by shoppers decreases, but the expected price paid by auto-renewers increases.*

Corollary 1 shows that when the PCW sets a higher fee, the expected price paid by both shoppers and auto-renewers rise. The fee is passed on by firms to consumers through a first-order stochastic shift in the prices set in equilibrium. Upon winning, a firm must pay the PCW for all of the $(\frac{n-1}{n})$ shoppers who purchase through the site.

Figure 1: Equilibrium price distributions calibrated with $q = 2$, $v = 1$, $\alpha = .7$

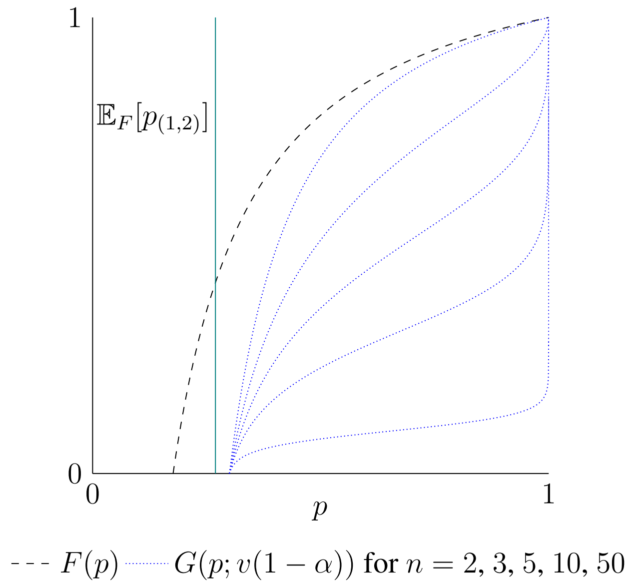

Because the amount paid rises with the fee, the price charged in equilibrium also rises with the fee.

The second effect is that the PCW increases competitive pressure among firms to fight for all $n - 1$ rival firms' shoppers. In contrast, without the PCW, firms effectively competed against only $q - 1$ rivals. Different models in the clearing house literature offer different predictions about the effect of n on equilibrium prices. Some derive distributions for which an increase in n raises prices for both types of consumer (e.g., Rosenthal, 1980); while in other models it lowers prices for shoppers and raises prices for captive consumers (e.g., Varian, 1980; Morgan et al., 2006).[17] My model belongs to this second category. An increase in n has two effects on equilibrium prices. First, increased competition pushes probability mass to the high-price extreme of the distribution, as Figure 1 shows. This results in a first-order stochastic ordering in n, so expected price increases in n and auto-renewers pay more. Second, shoppers now pay the lowest of $n + 1$ prices rather than n, which reduces the expected lowest price. Corollary 2 shows that this second effect more than offsets the first, implying that shoppers pay less

[17] Also see Baye et al. (2006) for a discussion on this point and a wider review of clearing-house models.

in expectation when there are more firms.[18]

In order to relate these two effects of a PCW to Proposition 3, I make the following remark.

Remark 1. *If* $q = n$, $G(p; 0) = F(p)$

The entrance of a PCW increases both the fee firms pay (from 0 to $v(1-\alpha)$) and the number of prices that shoppers compare (from q to n). The first effect raises the expected price auto-renewers pay, which is compounded further by $n > q$, so they are unambiguously worse off with a PCW. The effects pull shopper welfare in opposite directions, but Proposition 3 shows that no matter how large n is, it fails to undo the effect of the optimally chosen PCW fee c. Hence, shoppers are also worse off in expectation with a PCW.

Due to the constant-sum nature of the game, welfare necessarily sums to v in equilibrium. As firms make the same expected profit in both worlds, there is a one-to-one relation between the decrease in consumer welfare and the profits of the PCW. That the PCW does not reduce firm profits comes from the fact that (with or without the PCW) firms have $\left(\frac{1-\alpha}{n}\right)$ auto-renewers. Firms can therefore guarantee themselves $v\left(\frac{1-\alpha}{n}\right)$ in both worlds. Although this article focuses on consumer rather than producer welfare, it is important to point out that the incentives of the PCW and firms are not aligned. This is because there is exactly one cheapest price and hence $\alpha\left(\frac{n-1}{n}\right)$ shoppers who switch from non-cheapest prices, which increases with n. An increase in n however, would of course, squeeze per firm profit. Therefore, a PCW would always encourage market entry if it could.

One channel through which all consumers would gain is by more auto-renewers becoming shoppers:[19]

[18] Although the result of Corollary 2 is common to clearing-house models, it may seem a rather nuanced prediction of equilibrium pricing. Morgan et al. (2006) conduct an experiment with participants playing the role of firms against computerized buyers and found that when n was increased, prices paid by inelastic consumers indeed increased whereas those paid by shoppers decreased.

[19] This is a prediction common to clearing-house models, for which Morgan et al. (2006) find experimental support.

Corollary 3. *When pricing is determined by $G(p;c)$, as the proportion of shoppers α increases, expected prices for both types of consumer decrease.*

One may conjecture that a PCW would want to maximize α (the number of shoppers) in order to obtain more referral fees. However, this logic is incomplete. Expanding the PCWs action set to include the determination of α (one can think of the PCW determining α through advertising) yields the following result:[20]

Corollary 4. *If the PCW can determine α as well as c in the preliminary stage, then the PCW sets $\alpha = \frac{1}{2}$.*

PCW revenue is hump-shaped in profit. As $\alpha \to 0$ it receives less and less traffic, and hence vanishing revenue. As $\alpha \to 1$, firms have fewer auto-renewers to exploit, which pushes $v(1-\alpha)$, the maximum fee for which firms are willing to mix, to zero. Indeed, when $\alpha = 1$ all consumers are shoppers and the PCW removes any incentive for firms to increase prices as there are no auto-renewers to exploit. This leads any equilibrium to have all firms charging the same price, and such pure-strategy equilibria leave the PCW with zero profit. So even if the PCW could bring all consumers online, it has a strict incentive to ensure an interior solution for α.

5. Competing Aggregators

Suppose that there are now $K > 1$ PCWs indexed $k = 1, ..., K$ which move simultaneously in the first period each setting a fee c_k. Here I show that allowing competition between PCWs is not meaningful, per se. A crucial measure of competitive pressure is the number of aggregators shoppers check, r, where $1 \leq r \leq K$. In the setting without PCWs, shoppers know the prices of $q \leq n$ firms. In the world with PCWs, one can interpret $r \leq K$ as the number of aggregators that shoppers are aware of. In the equilibria derived, firms list on all PCWs so shoppers are indifferent to which set of PCWs they check. However, PCW fees and firm-pricing strategies depend on r. Hence there are

[20] α is determined costlessly for the PCW. If this advertising costs were a convex function of α, it would not change the results qualitatively.

different sets of equilibria for each r. This leads me to index equilibria by r. I first look at the case of $r = 1$, where shoppers check just one of the PCWs.

Proposition 4. *With $K > 1$ competing PCWs, if shoppers check $r = 1$ PCW then both types of consumer are worse off with the PCWs than without.*

Proposition 4 obtains because the high-fee, single-PCW equilibrium of Proposition 2 with $c = v(1 - \alpha)$ remains the unique equilibrium fee.[21] Introducing competing PCWs exerts no downward pressure on fees. There are no equilibria at lower fee levels because shoppers only check one PCW, which means there is no incentive for PCWs to undercut each other's fees. If they did, they would not increase their volume of sales, but they would receive lower fees. In contrast, for any candidate equilibrium with $c < v(1 - \alpha)$, there is an incentive to *raise* the fee because PCWs can maintain the volume of sales and earn a higher fee.

In the simplest textbook Bertrand result, there is an immediate fall in the equilibrium price between that of a monopoly firm and the marginal-cost pricing of two firms. At the aggregator level however, one can see that this logic is incomplete as Proposition 4 shows, with $K > 1$ such undercutting does not even necessarily begin. The persistence of the consumer-welfare-decreasing equilibrium in Proposition 4 is an artefact of each shopper only checking one PCW which effectively makes each PCW a monopoly. Because all PCWs offer the same information, shoppers have no incentive to check other PCWs. In this respect, the equilibrium is reminiscent of Diamond (1971), but at the aggregation not the firm level. One might think that the Bertrand remedy for shoppers would be to require them to check at least two PCWs, causing PCWs to undercut each other until an equilibrium with all PCWs charging $c = 0$ is reached. I now explain how this logic is incomplete. Suppose that the model is as before with $K > 1$ aggregators, but shoppers now check $r > 1$ PCWs. First, I examine the special case of $r = K > 1$:

Proposition 5. *With $K > 1$ competing PCWs, shoppers are guaranteed to be better off than before the introduction of PCWs if and only if they check all PCWs: The unique*

[21] As in the monopolist case, firms' outside option gives rise to pure-pricing firm responses at higher fee levels, which bounds equilibrium fee levels from above at $c = v(1 - \alpha)$; see the Appendix for details.

equilibrium PCW fee-level is $c = 0$ if and only if $r = K > 1$.

The spirit of this result resembles that of Bertrand. To see why there cannot be some other equilibrium with $c > 0$, suppose so and consider an undercutting deviation by PCW$_1$ to some $\hat{c}_1 = c - \epsilon$. Shoppers do not detect the deviation and so do not change their behavior. As for firms, notice that for any p they strictly prefer to list exclusively on PCW$_1$: When a firm is the cheapest, it will sell to *all* shoppers so long as it lists on *some* PCW. This is precisely because $r = K$. Hence by listing on PCW$_1$ only, there is no reduction in the number of shoppers switching to them when they are cheapest, but there is a reduction in the fee the firm pays as $\hat{c}_1 < c$. The PCW finds this deviation strictly profitable because it receives a discrete gain in the number of shoppers switching through it, for an arbitrarily small loss in price.

Where shoppers check more than one PCW, but not necessarily all PCWs, we have:

Lemma 2. *When $K > r > 1$, there exists an equilibrium in which PCWs charge $\bar{c} > 0$, firms list on all PCWs, mixing over prices by $G(p; \bar{c})$ where*

$$\bar{c} = \frac{v(1-\alpha)Kr(K-r)}{K(1+r(K-2)) + \alpha r(K-1)(r-1)(n-1)}.$$

More generally, there exist equilibria in which PCWs charge $c \in [0, \bar{c}]$. To understand how much consumer welfare can be reduced, I discuss the highest-fee equilibrium from this set. Because PCWs have a strong incentive to coordinate on this equilibrium, it may be especially relevant in practice.

Notice that substituting $r = K$ in Lemma 2 yields $\bar{c} = 0$, and one obtains Proposition 5. One can see now that the Bertrand-style reasoning underlying Proposition 5 was a special case. To see why the principle does not apply more generally, consider a fee level $c > 0$ and an undercutting deviation by PCW$_1$ to some $c_1 = c - \epsilon > 0$. Unlike when $K = r > 1$, when $K > r > 1$ it is not necessarily better for a firm to only list on the cheaper PCW$_1$. By listing a price exclusively on PCW$_1$, there are now $\frac{K-r}{K} > 0$ shoppers who do not see the firm's price, and so will not buy from it even if it is the cheapest. Firms now face a trade-off: Exclusively listing on PCW$_1$ means that any sales incur only the lower fee $\hat{c}_1$ upon a sale, but there will be a reduction in sales volume because not all shoppers check PCW_1. Which force is stronger in this trade-off depends on the size of the undercut ϵ. If PCW_1 undercuts by enough, firms will deviate to exclusively list on PCW_1, breaking the symmetric equilibrium. Unlike the simpler logic of Proposition 5, it is no longer true that any $\epsilon > 0$ undercut will attract firms to exclusively list on the cheapest PCW. Hence, PCWs do not always have an incentive to undercut each other and higher-price equilibria are sustained.

I now discuss how the set of equilibria varies with how many PCWs shoppers check (r) and the number of aggregators (K). Firstly, $\bar{c}$ is limited by a higher r: $\frac{d\bar{c}}{dr} < 0$. That is, as shoppers check more PCWs in equilibrium, the incentive for a PCW to undercut the fees of other PCWs increases so that only lower fee-levels can be sustained in equilibrium.

However, as the number of aggregators increases, the incentive is reversed. The number of shoppers checking a given PCW ($\frac{r}{K}$) falls. Accordingly, in equilibrium each firm receives less of its expected revenue from any single PCW. It then requires a more severe undercut from a PCW to get firms to exclusively list on it and forgo the business available from the other aggregators. For undercuts that are too severe, it is unprofitable for a PCW to deviate, even if it were to win exclusive arrangements with all firms as a result. Thus, as K increases, higher equilibrium fees can be sustained in equilibrium: $\frac{d\bar{c}}{dK} > 0$. This allows for the result that a higher number of aggregators can lead to higher fees, and hence higher prices. Furthermore:

Proposition 6. *For any r, there exists a $\tilde{K}$ such that as long as there are more than $\tilde{K}$ aggregators both types of consumer can be worse off than before the introduction of any PCWs.*

In the limit, the proportion of firm income that comes from sales on any one aggregator becomes vanishingly small. As this happens, $\bar{c} \to v(1-\alpha)$ i.e., sustainable equilibrium fee levels approach the monopoly-PCW level, again making both types of consumer worse off than before the introduction of the sites.

6. Publicly-Observable Fees

One reason that competing aggregators do not drive fees to zero is that shoppers do not detect changes in the fees of PCWs in equilibrium. This precludes a coordinated response between firms and shoppers that could punish a PCW that charges higher fees. If fees were publicly announced so that shoppers were aware of them, then credible subgame equilibria could follow the fee-setting decision in which the PCW charging the lowest fee is attended by all firms and shoppers. Any higher-fee equilibrium would then be undercut until $c = 0$, leading to lower shopper-prices. Indeed:

Proposition 7. *When $K > 1$ and PCW fees are observed by shoppers, there exists an equilibrium with $c = 0$.*

There are multiple equilibria because there are many subgame equilibria that can follow each value of the vector of PCW fees, including somewhat unintuitive ones where coordination occurs at more expensive PCWs.[22] If one adopts the plausible refinement that firms and consumers only patronize the lowest-fee PCWs, then one obtains the zero-fee equilibrium as the unique equilibrium.

However, even with sufficient refinement criteria to implement the zero-fee equilibrium, in some markets one may question a policy of fee-announcements on a more fundamental level. If fees can be publicly announced, then surely so can firm prices, which would extinguish the role of a PCW in the first place.

In reality, PCW fees are not publicized directly. However, some PCWs do advertise summary statistics of the price information of firms that list on them. PCWs frequently advertise the average savings a consumer using their site is expected to make, which could direct shoppers to the cheapest PCW. By Proposition 7, this could lead to a shopper-welfare-improving equilibrium. However, many PCWs do not advertise based on purchase-relevant information.[23] PCWs spend large sums on such advertising, which has been shown to correlate with the number of unique visitors they experience,

[22] Proposition 7 excludes the monopoly-PCW case ($K = 1$), where the unique equilibrium is still that of Proposition 2 because of course no coordination between firms and consumers over which PCW to attend is possible when there is only one PCW.

[23] See the campaign of Comparethemarket.com, based on a story about meerkats.

suggesting that many shoppers are directed to PCWs based on information other than price.[24] If such 'persuasive' advertising caused all shoppers to loyally visit one site each i.e., $r = 1$, Proposition 2 applies and both types of consumers would still have been better off without the PCW industry.

7. Price Discrimination

So far, I have considered the impact of web services that list or 'aggregate' the available information (prices) charged by firms offering a product or service. In practice, this is often the case in the markets for gas, electricity, financial products such as mortgages, and durable goods.[25] However, in other markets, a firm may set a price p_0 for a direct purchase, and p_k different prices for each PCW$_k$ that it lists on. Where a PCW operates by referring shoppers back to a firm's website to complete the purchase, the fact that the click came from a PCW is recognized by the firm's site, which then offers the price seen on the PCW that attracted the click. When $K \geq r = 1$, we have:

Proposition 8. *With price discrimination, if $r = 1$, there exists an equilibrium in which PCWs set $c = v(1 - \alpha)$, firms list on all PCWs, $p_0 = v$ and $p_1 = \cdots = p_K = v(1 - \alpha)$.*

The ability of firms to price discriminate does not prevent PCWs from setting fees at the same high level as in Proposition 4 because $r = 1$. Thus, as before, there is effectively no competitive pressure between PCWs. However, price discrimination does lead to firms listing a common price on PCWs in equilibrium because PCWs no longer have an incentive to keep prices posted on it dispersed. This is because firms can set a high direct purchase price $p_0 = v$ and a lower price through the PCWs. Given this, shoppers always purchase through a PCW. Because all α shoppers now purchase through PCWs, rather than in the case without discrimination, where only $\alpha\left(\frac{n-1}{n}\right)$ did so, total PCW profit is now $\alpha v(1 - \alpha)$ rather than $\left(\frac{n-1}{n}\right)\alpha v(1 - \alpha)$. As firm profit is unaffected, it follows then that relative to the case without price discrimination, total consumer

[24]The big four PCWs in the UK spent approximately £110m in 2010 on advertising. The evidence here is from Nielsen Company (findings reported by the This is Money 2015 and Marketing Magazine 2011).

[25]In the example of the UK gas and electricity markets, regulation has enforced that each energy company may only offer a maximum of four tariffs in total.

welfare falls. There are, however, opposing effects on shoppers and auto-renewers. Auto-renewers are worse off, as now firms charge them the monopoly price v. Shopper welfare though improves as they now face $p_k = v(1-\alpha)$ for sure, whereas before this was just the minimum of the support of equilibrium prices. More importantly, shopper welfare does not improve sufficiently to overturn Proposition 3, which continues to hold for this equilibrium, leaving both types are worse off than in a world with no PCWs. The ability of firms to discriminate allows them to fully extract surplus from their captive auto-renewers; but PCWs can now extract the revenue from sales to all shoppers through their sites.[26]

In equilibria with $r > 1$, the incentive for aggregators to undercut is present. Corollary 5 describes a best response of firms following a unilateral downward deviation by a PCW$_1$ from a symmetric fee level.

Corollary 5. *When firms price discriminate, following* $c_1 \in (0, c_2)$ *and* $c_2 = c_3 = , \ldots, = c_K \equiv c \in (0, v(1-\alpha)]$ *there exist mutual best-responses of firms such that they list on all PCWs, setting* $p_0 = v$ *and* $p_k = \cdots = c_k$ *for all* k.

When $r = 1$, there is no incentive for a PCW to make such an undercutting deviation as suggested by the equilibrium described in Proposition 8. When $r > 1$ however, PCWs have this incentive to undercut because they can enjoy a discrete gain in fee revenue from the $\frac{r-1}{K}$ of shoppers who were checking their PCW but buying though another site in the symmetric equilibrium. When firms can price discriminate across websites, firms can compete in prices on PCW$_1$ without changing their prices on other PCWs. This shows how price discrimination can unleash undercutting at the PCW level whenever consumers check $r > 1$ PCWs, which can in turn lead to zero-fee equilibria.[27] This contrasts with markets where PCWs aggregate price-information, where Proposition 5 showed that this undercutting was only fully unlocked when $r = K$.

[26] As before, in the equilibrium of Proposition 8, PCWs cannot raise fees further, to say $c' > v(1-\alpha)$ because of firms' outside option. Following such a unilateral PCW deviation, firms would set $p_0 = v(1-\alpha)$ so that their shoppers purchase directly from them, reducing PCW profit to zero.

[27] The existence of zero-fee equilibria when $r > 1$ are shown in the Appendix.

Aggregation and Discrimination in Large Markets

In the setting with a PCW and price discrimination, the equilibrium of Proposition 8 shows the price paid by the two types are as maximally separated: Shopper price is competed down to firms' marginal cost c, and auto-renewer price is v. In the setting with an aggregator and no price discrimination, the equilibrium is given by Proposition 2 where Corollary 2 explains that as the number of firms increases, the expected prices paid by shoppers and that paid by auto-renewers, diverge. This occurs because as the number of firms increases, so does the competitive pressure on pricing to win shoppers. As a result, more probability mass is placed on lower prices. Firms compensate for this by also increasing the mass placed on higher prices, increasing their expected profit from auto-renewers. In fact, for arbitrarily large markets (where $n \to \infty$), I show that these two settings are equivalent.

Proposition 9. *For $n \to \infty$, following the introduction of a PCW, the expected prices faced by both types of consumer in the setting without price discrimination (Proposition 2) are the same as in the setting with price discrimination (Proposition 8).*

The result highlights the connection between the two market structures. Indeed, I found that both types of consumer can be worse off with a PCW with or without the possibility price discrimination. I emphasize therefore that the key difference between the settings lies in their predictions under competition at the aggregator level.

8. The Extensive Search Margin

This paper utilizes a clearing-house framework, where auto-renewers are inactive and can also be interpreted as being offline, loyal, uninformed or as having high search costs. Here, I focus on a search-cost rationalization, better applied to markets where obtaining a quote requires more information from the consumer e.g., home insurance. In an environment without a PCW where auto-renewers find it too costly to enter these details into a firm's website to retrieve *one* extra price, the introduction of a PCW offers to expose *all* prices to them, for the same search cost. Depending on their search cost, the introduction of a PCW may then cause an auto-renewer to engage with the market via the PCW. Some empirical studies have offered a similar argument to explain observed increases in market competitiveness (e.g., Brown and Goolsbee, 2002; Byrne et al., 2014). Their arguments are distinct from mine because they contrast a world with web-based aggregators relative to a world without the Internet, rather than a world with the Internet and firm websites. This engagement of new customers is commonly referred to as the 'extensive search margin' (for a recent discussion see Moraga-González et al., 2015).

The benefit of an additional search for a consumer in the world without a PCW is the difference between the expected price and the expected lowest of two prices drawn (from F). In the world with a PCW, the benefit of a search on the PCW is the difference between the expected price and the expected lowest of n draws (from G).[28] I denote these search benefits without and with an aggregator respectively as,

$$\mathcal{B}_0 = \mathbb{E}_F[p] - \mathbb{E}_F[p_{(1,2)}], \qquad \mathcal{B}_1 = \mathbb{E}_G[p] - \mathbb{E}_G[p_{(1,n)}].$$

The model is as before but now each auto-renewer faces a search cost s. I assume these costs are heterogeneous, distributed by S over $s \in [\underline{s}, \infty)$.[29] I assume $\underline{s} > \mathcal{B}_0$ which means that without a PCW, no auto-renewers shop, preserving the equilibrium of Proposition 1. After the introduction of a PCW, the benefit of a search ($\mathcal{B}_1$) may outweigh the cost (s) for some auto-renewers, who then choose to use the site. I use the term 'converts' and 'non-converts' to distinguish between auto-renewers who decide to shop or not in equilibrium.

The total number of consumers using the PCW (shoppers and converts) is endogenously determined in equilibrium and is denoted $\tilde{\alpha} = \alpha + (1-\alpha)S(\mathcal{B}_1)$. Given $\tilde{\alpha}$, the PCW sets its profit-maximizing fee $c = v(1-\tilde{\alpha})$. As c and $\tilde{\alpha}$ are exogenous to firms, equilibrium pricing is as in Proposition 2 with $\tilde{\alpha}$ replacing α. In turn, pricing determines $\mathcal{B}_1$. There is an equilibrium when this value of $\tilde{\alpha}$ satisfies $S(\mathcal{B}_1) = \frac{\tilde{\alpha}-\alpha}{1-\alpha}$.

[28] These terms are analogous to the 'value of information' in Varian (1980).

[29] That there is no upper bound ensures that there are always some auto-renewers, and hence price dispersion, in equilibrium.

When there exists such an S, $\tilde{\alpha}$ is said to be 'rationalized'.

Corollary 3 showed that a higher α increases the welfare of all consumers. Corollary 1 showed that a lower c has the same effect. As the equilibrium fee level is $v(1-\tilde{\alpha})$, both forces work to benefit all types of consumer. However, whether consumers actually gain depends on how many auto-renewers are converted. In fact, relative to the world without a PCW, the presence of converts is not sufficient to guarantee lower prices for any consumer, not even converts themselves:

Proposition 10. *Shoppers, converts and non-converts can all be worse off with a PCW than without.*

The proof gives an example of such an equilibrium, along with a distribution S that rationalizes it. More generally, there can be many equilibria, each with a different $\tilde{\alpha}$. If the benefit ($\mathcal{B}_1$) is small or there are no types with low search costs so that $S(\mathcal{B}_1) = 0$, there are no converts ($\tilde{\alpha} = \alpha$) and the equilibrium of Proposition 2 applies. Proposition 10 shows that when *some* auto-renewers convert, all consumers can be worse off with a PCW. However, there may also exist equilibria where $\tilde{\alpha}$ is high enough such that some consumers benefit. Whether these equilibria exist depends on the distribution of search costs S. For a given PCW search benefit B_1, when more auto-renewers have low search costs (higher $S(\mathcal{B}_1)$), $\tilde{\alpha}$ is higher, c is lower and total consumer welfare is higher.

The number of firms also determines the size of the benefit the PCW offers ($\mathcal{B}_1$), and hence the number of converts. I now investigate which consumers benefit from an aggregator when the potential benefit it offers to consumers ($\mathcal{B}_1 - \mathcal{B}_0$) is as large as possible. For a given $\tilde{\alpha}$, a higher number of firms increases the equilibrium search benefit $\mathcal{B}_1$ (see Corollary 2). Specifically, as $n \to \infty$, $\mathcal{B}_1 \to v - c$ (see Proposition 9), which is as large as possible. To maximize $B_1 - B_0$ I let $q = 2$ in the world without a PCW, which makes $\mathcal{B}_0$ as low as possible, as $\mathcal{B}_0$ is increasing in q.[30] Accordingly, I define,

Definition. *A market has 'maximum potential' when $q = 2$ and $n \to \infty$.*

[30]See the end of the proof of Proposition 5.

Proposition 11. *If the market has maximum potential: converts are better off; shoppers can be worse off; and non-converts are worse off with a PCW than without.*

When the market has maximum potential, converts are guaranteed to be better off but shoppers still may not be.[31] This is because there still may not be enough auto-renewers converting to PCW use.

9. Conclusion

The analysis shows that the introduction of PCWs may not in fact benefit consumers by reducing expected prices. The introduction of a single aggregator facilitates comparison of the whole marketplace for shoppers, which exerts competitive pressure on firm pricing. However, the aggregator charges a fee which, in turn, places upward pressure on prices. The net effect is that prices increase for all consumers, who would be better off without the site.

Competition at the aggregator level need not lead to a reduction in fees. There are many equilibria, which I parameterize by the number of PCWs shoppers check, or are informed of. If shoppers only check one PCW each, consumers are no better off than in the monopoly setting. A greater number of aggregators is only guaranteed benefit shoppers if they check all of them. If shoppers check an intermediate number, increasing the number of aggregators can lead to higher prices; for a sufficiently high number, all consumers can again be better off without the industry. Hence, when there are many aggregators in the market, how many of them shoppers check is a crucial variable. As a result, regulatory bodies may wish to consider incentivizing consumers to check more, alongside stronger actions such as limiting the fees charged by aggregators and limiting the number of PCWs in the market.

If competing PCWs publicly announce fees, this can result in low-fee equilibria, but it relies on coordination between firms and shoppers. In markets with price discrimination, if shoppers only check one PCW, the monopoly fee level can still be sustained,

[31] By Proposition 9, one can also use Proposition 11 to consider the effect of introducing a PCW into a market with search costs and price discrimination under the equilibrium of Proposition 8.

making consumers worse off. However, with competing aggregators where shoppers check multiple PCWs, there is then the incentive for the sites to undercut each other's fees. As such, regulators may also wish to consider whether it is possible to introduce price discrimination into markets in which it is not currently present. In online markets with non-negligible search costs, even those consumers who rationally start engaging in price comparison may be worse off following the introduction of a PCW. Helpful policies would help erode these costs where possible and encourage more inactive consumers to engage in comparison.

Appendix

A1. Summary Statistics of UK PCW Industry

The figures quoted for turnover and profit in the introduction for the UK utilities and services industry are estimates for 2013 for the largest four such companies. The estimate for each site was taken from the following sources:

- Money Supermarket: Their own 2013 annual report.
- Confused: The Admiral Group's Annual Reports 2012 and 2013.
- Go Compare: Newspaper 'The Guardian' 2014 article.
- Compare the Market: Parent BGL Group's Annual Report 2013 and BBC article 2015.

For the first three, the figures were taken directly from the cited sources. For Compare the Market, estimated to be the largest of the four sites, the estimate is particularly rough as BGL Group offer no breakdown of their accounts. I assumed that the proportion of BGL's total revenue and profit due to Compare the Market was the same where the estimate for annual profit due to Compare the Market is taken from the BBC article. Even if the estimated £800m of total turnover and 14% for average annual profit growth for these sites is off by a margin, the industry can still be considered large and growing.

A2. A World Without a PCW

The domain of prices is $\mathbb{R}$. The equilibrium pricing strategy can always be described by its CDF (denoted F in this section, and with other letters later on). In what follows, either equilibrium pricing distributions will be pure (so that F is flat, with one jump discontinuity at this price); or will have no point masses so that F is continuous, which implies the density f exists and $f = F'$, wherever F' exists.

Lemma A1. *In equilibrium, there are no prices p s.t. $p \leq 0$ or $p > v$.*

Proof: Any $p \leq 0$ generates firm profits $\pi(p) \leq 0$ which is dominated by $p = v$ which gives profit of at least $v\frac{(1-\alpha)}{n} > 0$ because firms always sell to their auto-renewers. For $p > v$, $\pi(p) = 0$ because no one will buy at such a high price, hence again $p = v$ dominates. ■

Lemma A2. *There are no point masses in equilibrium.*

Proof of Lemma A2: This is a variation of Proposition 3 in Varian (1980). Suppose not. Then the there is a point mass in equilibrium, $\dot{p}$ s.t. $pr(p = \dot{p}) > 0$. Note that $\dot{p} \in (0, v]$ from Lemma A1. Because the number of point masses must be countable, there exists an $\varepsilon > 0$ small such that $\dot{p} - \varepsilon > 0$ and is charged with probability zero. Consider a deviation of a firm from the equilibrium F to a distribution over prices where the only difference is that the new distribution charges $\dot{p} - \varepsilon$ with probability $pr(p = \dot{p})$ and $\dot{p}$ with probability zero. Note that a firm appears in $\binom{n-1}{q-1}$ of the groups of shoppers. Index these groups $z = 1, ..., \binom{n-1}{q-1}$ (this is without loss as F is symmetric). Let $pr(\dot{p}; t, z)$ be the probability under F that a firm is cheapest in group z along with $t = 0, ..., q-1$ others (due to symmetry, $pr(\dot{p}; t, z)$ is the same for all groups, so let this more simply be termed $pr(\dot{p}; t)$). Call the difference in profit due to the deviation d and note that:

$$\lim_{\varepsilon \to 0} d = \sum_{z=1}^{\binom{n-1}{q-1}} \sum_{t=0}^{q-1} pr(\dot{p}; t, z) \frac{\alpha}{\binom{n}{q}} \left(1 - \left(\frac{1}{q} + \frac{q-(t+1)}{q(t+1)}\right)\right) \dot{p}$$

The term in parentheses is the difference in the amount of shoppers won under the deviation and under F, in given group, when the firm along with t others charge $\dot{p}$, and $\dot{p}$ is the lowest price in that group. Due to symmetry of the groups and that $pr(\dot{p}; t) > 0$

(because $\dot{p}$ is a point mass in F), this simplifies to:

$$\frac{\binom{n-1}{q-1}}{\binom{n}{q}} \sum_{t=1}^{q-1} pr(\dot{p};t)\alpha\left(\frac{t}{t+1}\right)\dot{p} > 0$$

hence for some $\varepsilon > 0$ there exists a profitable deviation, so F could not have been an equilibrium. ■

Lemma A3. *In any equilibrium, the maximum of the support of f must be v.*

Proof: Suppose not. Define $\bar{p}$ as the maximum element of the support and note that by Lemma A2 the probability of a tie at any price is zero. By Lemma A1, $\bar{p} \in (0, v)$. Consider a firm called upon to play $\bar{p} < v$ in equilibrium. They only sell to their auto-renewers, making $\bar{p}\frac{(1-\alpha)}{n}$ but would strictly prefer to deviate to v and make $v\frac{(1-\alpha)}{n}$, a contradiction. ■

Proposition 1. *In the unique equilibrium firms mix according to the distribution*

$$F(p) = 1 - \left[\frac{(v-p)(1-\alpha)}{qp\alpha}\right]^{\frac{1}{q-1}} \text{ over the support } p \in \left[\frac{v(1-\alpha)}{1+\alpha(q-1)}, v\right].$$

Proof: By Lemma A2, there is more than one element of the equilibrium support, and by Lemma A3 v is the maximal element. In equilibrium, a firm must be indifferent between all elements of the support p, hence profit must equal $v\left(\frac{1-\alpha}{n}\right)$ for all p, that is:

$$v\left(\frac{1-\alpha}{n}\right) = p\left[\frac{1-\alpha}{n} + \frac{\alpha}{\binom{n}{q}}X(p)\right] \tag{A1}$$

$$X(p) \equiv \binom{n-1}{q-1}\binom{n-1}{n-1}F(p)^0(1-F(p))^{n-1} + ... + \binom{q-1}{q-1}\binom{n-1}{q-1}F(p)^{n-q}(1-F(p))^{q-1}$$

The first term on the RHS of (A1) is the profit from ARs, who always purchase at $p \leq v$. The second term is the expected proportion of shoppers that a firm will win, charging price p. Shoppers can be characterized by $\binom{n}{q}$ groups, where the set of groups is given by $\{1, ..., n\}^q$. $X(p)$ describes the expected number of groups a firm expects to win given it charges p. By Lemma A2 there are no ties in price, so label prices by $p_1 < ... < p_n$. p_1 will be the cheapest in every group in which it appears, and it appears in $\binom{n-1}{q-1}$ of the groups. The probability of being the lowest price is given by $\binom{n-1}{n-1}F(p)^0(1-F(p))^{n-1}$, which accounts for the first term in $X(p)$. The observation that p_i is the cheapest in $\binom{n-i}{q-1}$ groups if $i \leq n-(q-1)$, zero groups otherwise, accounts for the remaining terms of $X(p)$. The following manipulations to simplify the second term on RHS of (A1) this make use of the binomial theorem:

$$= p\frac{\alpha}{\binom{n}{q}}\sum_{j=q-1}^{n-1}\binom{j}{q-1}\binom{n-1}{j}F(p)^{n-j-1}(1-F(p))^{j}$$

which after some manipulations can be shown to be:

$$= p\alpha\frac{q}{n}(1-F(p))^{q-1}$$

rearranging for $F(p)$ gives:

$$F(p) = 1 - \left[\frac{(v-p)(1-\alpha)}{qp\alpha}\right]^{\frac{1}{q-1}}$$

Notice that this is a well-defined c.d.f. over:

$$p \in \left[\frac{v(1-\alpha)}{1+\alpha(q-1)}, v\right]$$

Notice that v is strictly preferred to any $p \in \left[0, \frac{v(1-\alpha)}{1+\alpha(q-1)}\right)$. ■

A3. A World With a PCW

I look for symmetric equilibria where PCWs charge some fee level $c \geq 0$ and shoppers check PCWs in equilibrium. I do not look at equilibria where firms never list on PCWs, where the setting without a PCW applies. To derive equilibria, one needs to know the mutual best-responses of firms to unilateral deviations of PCWs. To do so, take $\mathbf{c}$ and equilibrium shopper strategy as given, and consider the firm responses.

Let $K \geq 1$ denote the number of PCWs and $r : K \geq r \geq 1$ the number of PCWs checked by shoppers. In symmetric equilibrium, a proportion $\frac{r}{K}$ of each firm's shoppers check any given PCW. Define a vector of PCW fees as $\mathbf{c} = (c_1, ..., c_K) \in \mathbb{R}_+^K$ labeled such that $c_1 \leq ... \leq c_K$. Let $\beta_k \in [0, 1]$ be the probability with which a firm enters PCW$_k$ and define the event E: "all PCWs are empty". Denote $(a_1, a_2) = (p, \mathcal{K})$ as a firm's action, where p is the price charged and $2^{\{1,...,K\}}$ is set of all combinations of PCWs they could choose to list in where $\mathcal{K}$ a typical element, and $\emptyset$ denotes not listing on any PCW. Define the following CDF, which is used throughout:

$$G(p;\mathbf{c}) = 1 - \left[\frac{(v-p)(1-\alpha)}{\alpha\left(np-(n-1)\frac{1}{K}(c_1+\cdots+c_K)\right)}\right]^{\frac{1}{n-1}} \tag{A2}$$

which is well-defined over the support $[\underline{p}(\mathbf{c}), v]$ where $\underline{p}(\mathbf{c}) = \frac{v(1-\alpha)+\frac{1}{K}(c_1+\cdots+c_K)\alpha(n-1)}{1+\alpha(n-1)}$. When $c_1 = \cdots = c_K \equiv c$, let $G(p;\mathbf{c})$ and $\underline{p}(\mathbf{c})$ also be written $G(p;c)$ and $\underline{p}(c)$.

Lemma A4. *In any equilibrium, there are no prices p charged s.t. $p \leq 0$ or $p > v$.*
Proof: No $p \leq 0$ or $p > v$ because they yield negative and zero profit respectively, whereas $(v, \emptyset)$ yields $\frac{v(1-\alpha)}{n} > 0$. ■

Lemma A5. *If $c_1 \in [0, v(1-\alpha))$, $pr(E) = 0$.*
Proof: Suppose $pr(E) > 0$. Denote the infimum of prices charged when no PCW is listed on and that of prices ever listed on a PCW as p_0 and $\underline{p}$ respectively. [Note the infima exist because prices are a bounded from below by Lemma A4]. Note that $p_0 \geq \underline{p}$ because $(\underline{p}, \emptyset)$ is strictly preferred to any lower price as a firm faces no competition for prices below $\underline{p}$ off the PCWs. Consider when the firm is called upon to play $(p_0, \emptyset)$ (or a price arbitrarily closely above p_0):

If $p_0 > c_1$, a deviation to $(p_0 - \epsilon, 1)$ is strictly profitable. This is because with probability $pr(E)^{n-1} > 0$ PCWs are empty with other firms charging at least p_0. By listing the firm then has a positive probability of winning $\alpha\frac{r}{K}\frac{n-1}{n}$ new shoppers. For a sufficiently small $\epsilon > 0$, this will offset the arbitrary loss in revenue from its own consumers.

If $p_0 \leq c_1$, firm profit must be at least $\frac{\underline{p}}{n}$ which can be guaranteed by $(\underline{p}, 1)$, because $p_0 \geq \underline{p}$. In turn, this must be at least as much as $\frac{v(1-\alpha)}{n}$ which the firm can guarantee by playing $(v, \emptyset)$. Putting these together, $p_0 \geq \underline{p} \geq v(1-\alpha) > c_1$ which contradicts $p_0 \leq c_1$. ■

Lemma A6. *If $c_1 \in [0, v(1-\alpha))$, firm strategies have no point masses.*
Proof: Define $\underline{p}_1(c_1)$ as the price at which a firm is indifferent between selling to all shoppers exclusively through the cheapest PCW(s) and charging v, only sell to auto-renewers:

$$\underline{p}_1(c_1) = \frac{v(1-\alpha)+\alpha c_1(n-1)}{1+\alpha(n-1)}$$

By Lemma A5, some $(p, \mathcal{K})$ is played. Note that firms would by construction not play $(p, \mathcal{K})$ where $p < \underline{p}_1(c_1)$, strictly preferring $(v, \emptyset)$. To see that there are no point masses:

If there were a point mass at $(\underline{p}_1(c_1), \mathcal{K})$ then there is a positive probability of being tied for the lowest price at $(\underline{p}_1(c_1), \mathcal{K})$. By definition of $\underline{p}_1(c_1)$ firms would strictly prefer to deviate to $(v, \emptyset)$.

If there were a point mass at $(\dot{p}, \mathcal{K})$ s.t. $\dot{p} > \underline{p}_1(c_1)$ then there is a positive probability of being tied for the lowest price at $(\dot{p}, \mathcal{K})$. A firm would strictly prefer to shift that probability mass to $(\dot{p} - \epsilon, \mathcal{K}')$ where $\mathcal{K}' = \mathcal{K} \setminus \{k : c_k \geq \dot{p}\}$. Here, the firm would sell to $\alpha \frac{r}{K} \frac{n-1}{n}$ other firms' shoppers at an arbitrary loss in revenue from its own consumers. There is always an $\epsilon > 0$ small enough to ensure this is profitable because $\underline{p}_1(c_1) > c_1 \iff c_1 < v(1 - \alpha)$. ■

Lemma A7. *If there are no point masses, the maximum of the support f must be v.*

Proof: This is a variant of Varian (1980) Proposition 7. ■

Lemma A8. *If $c_1 \in [0, v(1 - \alpha))$ and $c_1 < c_2$, $\beta_1 = 1$.*

Proof: By Lemma A5 it is never the case that all PCWs are empty. Suppose $\beta_1 < 1$. Lemma A6 implies there is more than one price, $\hat{p}$, that is listed on some other PCWs. By Lemma A7 there is one such that $\hat{p} < v$. Consider a firm being called upon to play this $(\hat{p}, \mathcal{K}_{\text{-}1})$ where $1 \notin \mathcal{K}_{\text{-}1}$ and $\bar{m} = \max\{\mathcal{K}_{\text{-}1}\}$. As this price has a positive probability of being the lowest of all firms, it will generate sales through the PCWs in $\mathcal{K}_{\text{-}1}$. But as PCW$_1$ is the unique cheapest PCW, there is a strictly profitable deviation to $(\hat{p}, \mathcal{K}_{\text{-}1} \cup 1 \setminus \bar{m})$. ■

Lemma A9. *If $c_1 = \ldots, c_K \equiv c \in (v(1 - \alpha), v]$, firm pricing strategies are pure where $p \in [v(1 - \alpha), c]$. Any $(\beta_1, \ldots, \beta_K) \in (0, 1]^K$ can be supported.*

Proof: Either there is a point mass or there is not.

1. Suppose there is a point mass at price $\dot{p}$. If $\dot{p} > c$, firms have a strict incentive to shift this mass to $(\dot{p} - \epsilon, \{1, \ldots, K\})$. If $\dot{p} < v(1 - \alpha)$, firms prefer to shift this mass to $(v, \emptyset)$. These leaves $\dot{p} \in [v(1 - \alpha), c]$ as the only points that can be point masses. There can be at most one point mass: If not, then there a second

point mass $\ddot{p} < c$, which if played with $\mathcal{K} \neq \emptyset$ would generate negative profit, so $(\ddot{p}, \emptyset)$ is preferred; if $\mathcal{K} = \emptyset$, then $(\ddot{p} + \epsilon, \emptyset)$ for some sufficiently small $\epsilon > 0$ is preferred. To see that this pure pricing at $p \in [v(1-\alpha), c]$ can be part of firm strategies, note that firm profit is $\pi = \frac{p}{n} \geq \frac{v(1-\alpha)}{n}$, so there is no strict incentive to sell only to auto-renewers instead. Because shoppers buy directly when prices are all the same, there are no sales through PCWs and so firms are indifferent between any $(\beta_1, \ldots, \beta_K) \in (0,1]^K$.

2. Suppose there is no point mass. By Lemma A7 the maximum of the support is v, where v is not the only element of the support, else it would be a point mass. There is therefore a positive probability of a firm being the cheapest at some p. There can be no $(p, \mathcal{K})$ s.t. $p < c$ charged: If $\mathcal{K} \neq \emptyset$ played profit from these sales is negative, so $(p, \emptyset)$ is preferred; if $\mathcal{K} = \emptyset$, then $(p + \epsilon, \emptyset)$ for some sufficiently small $\epsilon > 0$ is preferred. Given $p \in [c, v]$, $pr(E) = 0$. This follows because firms strictly prefer $(p, \{1, \ldots, K\})$ to $(p, \emptyset)$ for all $p \in (c, v)$. For any $1 \leq r \leq K$ firms are content to list prices on at least as many PCWs as is necessary to make sure every shopper sees their price e.g., for $r = K = 1$ all of them; for $r = K$, just one of them. There can therefore, be different configurations of β_K's depending on r, K so long as $pr(E) = 0$. To determine firm pricing strategy it must be that firms are indifferent between every p they are called upon to play:

$$v\left(\frac{1-\alpha}{n}\right) = p\left(\frac{1-\alpha}{n}\right) + (1 - G(p;c))^{n-1}\left[\frac{\alpha}{n}p + \frac{\alpha(n-1)}{n}(p-c)\right]$$

which can be re-arranged to give $G(p;c)$ from (A2). However, $\underline{p}(c) < c$ because $c > v(1-\alpha)$, so firms would make strictly negative profits at prices $p \in (\underline{p}(c), c)$, preferring not to list. This provides a contradiction, so there do not exist strategic firm responses with no point masses.

■

A4. Results for $\mathbf{K = 1}$ or $\mathbf{r = 1}$

Lemma A10. *If $r = 1$, $c_1 = \cdots = c_{K-1} \in [0, v(1-\alpha))$ and $c_K \in [c_1, \underline{p}(\mathbf{c}))$: $\beta_k = 1$ for all k where firms mix by CDF $G(p; \mathbf{c})$.*

Proof: By Lemma A5 there is always some price posted on some PCW(s). By Lemma A6 these prices are not point masses and by Lemma A7 the maximum of the support is v, where only auto-renewers are sold to. There is therefore a positive probability of a sale through some PCW(s) at some $(p, \mathcal{K})$. If $k \in \mathcal{K}$ where $k < K$ and sales there are profitable (when $p > c_1$), then $\{1, \ldots, K-1\} \in \mathcal{K}$; if $K \in \mathcal{K}$ and sales there are profitable $p > c_K$ (whether or not $p > c_K$ is the only consideration because $r = 1$), it follows that $\{1, \ldots, K\} \in \mathcal{K}$. To be part of firm strategy, it must also be that firms prefer to play $(p, \mathcal{K})$ than to charge v and sell only to their auto-renewers. Note that $\underline{p}(\mathbf{c})$ from (A2) is the price at which firms are indifferent between selling through $\{1, \ldots, K\}$ with certainty and selling only to their auto-renewers. Similarly, denote $\underline{p}_{\text{-}K}(\mathbf{c})$ as the indifference point between selling for sure on $\{1, \ldots, K-1\}$ and charging v, and note $\underline{p}_{\text{-}K}(\mathbf{c}) < \underline{p}(\mathbf{c})$. There will then be no $(p, \{1, \ldots, K-1\})$ played s.t. $p < \underline{p}_{\text{-}K}(\mathbf{c})$ and no $(p, \{1, \ldots, K\})$ s.t. $p < \underline{p}(\mathbf{c})$. For prices close to v, $\mathcal{K} = \{1, \ldots, K\}$. To determine firm pricing strategy, it must be that firms are indifferent between every $(p, \{1, \ldots, K\})$ they are called upon to play:

$$v\left(\frac{1-\alpha}{n}\right) = p\left(\frac{1-\alpha}{n}\right) + (1-G(p;\mathbf{c}))^{n-1}\left[\frac{\alpha}{n}p + \frac{\alpha(n-1)}{n}\left(p - \left(c_1\frac{K-1}{K} + c_K\frac{1}{K}\right)\right)\right]$$

which can be re-arranged for $G(p; \mathbf{c})$ to give the CDF from (A2). Because the minimum of the support is $\underline{p}(\mathbf{c}) > c_K \geq c_1$ (the first relation follows because $c_1 < v(1-\alpha)$), all prices in the support generate profitable sales through all the PCWs, there is no price charged s.t. $p \in [\underline{p}_{\text{-}K}(\mathbf{c}), \underline{p}(\mathbf{c}))$. It follows that $\beta_k = 1$ for all k. For $K = 1$, let $c_1 \in [0, v(1-\alpha))$ and set $K = 1$ in the expressions of the Lemma. ■

Lemma A11. *If $r = 1$ and $c = v(1-\alpha)$ there are the following firm responses:*

1. *Pure-price strategies where $p = v(1-\alpha)$ is the only price ever charged. Here, any $(\beta_1, \ldots, \beta_K) \in (0, 1]^K$ can be supported.*

2. *Mixed-price strategies where $\beta_k = 1$ for all k and prices are distributed according to the CDF $G(p; v(1-\alpha))$ where $\underline{p}(v(1-\alpha)) = v(1-\alpha)$.*

Proof: Either there is a point mass or there is not.

1. Suppose there is a point mass at price $\dot{p}$. If $\dot{p} > c$, firms have a strict incentive to shift this mass to $(\dot{p} - \epsilon, \{1, \ldots, K\})$. If $\dot{p} < c$, firms have a strict incentive to

shift this mass to $(v, \emptyset)$. These leaves $\dot{p} = c$ as the only point that can be a point mass. To see that this pure pricing can be part of firm strategies, note that firm profit is $\pi = \frac{c_1}{n} = \frac{v(1-\alpha)}{n}$, so there is no strict incentive to sell only to auto-renewers instead. Because shoppers buy directly when prices are all the same, there are no sales through PCWs and so firms are indifferent between any $(\beta_1, \dots, \beta_K) \in (0, 1]^K$.

2. Suppose there is no point mass. By Lemma A6, v is the maximum of the support of prices. No prices $p < c$ are charged because $(v, \emptyset)$ is strictly preferred. For $p > c$, sales through all PCWs are profitable so $\beta_k = 1$ for all k. When v is played, firm profit is $\pi(v) = \frac{v(1-\alpha)}{n}$. To determine firm pricing strategy, it must be that firms are indifferent between every p they are called upon to play:

$$\pi(v) = p\left(\tfrac{1-\alpha}{n}\right) + \tfrac{\alpha}{n}(1 - G(p; v(1-\alpha)))^{n-1}\left[np - (n-1)v(1-\alpha)\right]. \qquad \text{(A3)}$$

which can be re-arranged to give the CDF from (A2).

■

Lemma A12. *If $r = 1$, $c_1 = \cdots = c_{K-1} = v(1-\alpha)$ and $c_K \in (c_1, v]$ there exist firm responses in pure-price strategies where $v(1-\alpha)$ is the only price ever charged. Here, any $(\beta_1, \dots, \beta_K) \in (0, 1]^K$ can be supported.*

Proof: In such an equilibrium, firm profit is $\pi = \frac{v(1-\alpha)}{n}$. Consider a deviation to p. If $p \in (v(1-\alpha), v]$, deviation profit is $\hat{\pi} = \frac{p(1-\alpha)}{n} \leq \pi$. If $p < v(1-\alpha)$, firms have a strict incentive to shift this mass to $(v, \emptyset)$. Because shoppers buy directly when prices are all the same, there are no sales through PCWs and so firms are indifferent between any $(\beta_1, \dots, \beta_K) \in (0, 1]^K$. ■

Lemma 1. *The mutual best-responses of firms as a function of c:*

1. *If $c \in [0, v(1-\alpha))$, firm responses have no point masses, and are described by $G(p; c)$.*

2. *If $c = v(1-\alpha)$, there are two classes of response, one with no point masses described by $G(p, c)$, and those that are degenerate.*

3. *If* $c \in (v(1-\alpha), v]$*, firm responses are degenerate.*

Where pricing is described by $G(p;c)$ *each firm always lists its price on the PCW.*

Proof: See Lemmas A10, A11 and A9 respectively. ■

Proposition 2. *In the unique equilibrium where shoppers check the PCW, the PCW sets a click-through fee of* $c = v\,(1-\alpha)$*, firms list on the PCW and mix over prices according to* $G(p; v(1-\alpha))$ *over the support* $p \in [v(1-\alpha), v]$.

Proof:

If $c \in [0, v(1-\alpha))$, Lemma A10 shows that there is a profitable upward deviation to $c \in (c_1, \underline{p}(c))$. It is profitable because there is an increase in fee-level but no reduction in the quantity of sales.

If $c \in (v(1-\alpha), v]$, Lemma A9 shows that firms will play pure-price strategies and so $u = 0$. If $c = v(1-\alpha)$ and firms play pure-pricing strategies, then $u = 0$ again. In these cases, Lemma A10 shows that a deviation to $c \in (0, v(1-\alpha))$ will generate $u_1 > 0$, so there are no equilibria where $c \in (v(1-\alpha), v]$ or for $c = v(1-\alpha)$ when firms respond with pure-pricing strategies.

If $c = v(1-\alpha)$ and firms mix over prices by $G(p; v(1-\alpha))$ as in Lemma A11, equilibrium PCW profit is $u^* = v(1-\alpha)\frac{\alpha}{K}\frac{n-1}{n} > 0$. Lemma A9 shows that any upward deviation would yield $u = 0$. There can be no profitable deviation downwards because, as Lemma A10 shows, the fee would be reduced for no gain in the quantity of sales. ■

Proposition 3. *Both types of consumer are worse off with the PCW than without.*

Proof of Proposition 3: First I show that auto-renewers are worse off under $G(p; v(1-\alpha))$ than $F(p)$ (referred to as G and F here). Auto-renewers pay the price quoted by their current firm. To show they are worse off with the PCW, I show that $\mathbb{E}_F[p] < \mathbb{E}_G[p]$ by showing that G first-order stochastic dominates (FOSDs) F. The distributions share the same upper bound on their supports, with F having a lower lower bound. Hence, G FOSDs F if $G(p; v(1-\alpha)) \leq F(p)$ for $p \in [v(1-\alpha), v]$, which can be re-arranged as

$$\left[\frac{(v-p)(1-\alpha)}{\alpha q p}\right]^{\frac{1}{q-1}} \leq \left[\frac{(v-p)(1-\alpha)}{\alpha n p - \alpha v(1-\alpha)(n-1)}\right]^{\frac{1}{n-1}}.$$

This holds because the terms in parentheses are in [0,1], the power on of the LHS is larger and the denominator on the LHS is larger iff $p(n-q) \leq v(1-\alpha)(n-1)$ which is satisfied because $p \leq v(1-\alpha)$ and $n-q < n-1$.

To show shoppers are worse off under G than F, first show that $\mathbb{E}_F[p_{(1,2)}]$ is lower than the lower bound of the support of g in the case of $q = 2$. Then, I use Proposition 3 of Morgan et al. (2006) which corresponds to my setup (the only difference is that they have $v = 1$), which states that $\mathbb{E}_F[p_{(1,2)}]$ is decreasing in q. Hence I show the first step here to prove that $\mathbb{E}_F[p_{(1,2)}]$ is below $v(1-\alpha)$ for all q. For $q = 2$,

$$\mathbb{E}_F\left[p_{(1,2)}\right] = \int\limits_{v\left[\frac{1-\alpha}{1+\alpha}\right]}^{v} f\left(p_{(1,2)}\right) p dp$$

where $f\left(p_{(1,2)}\right) = 2\left(1-F(p)\right) f(p)$ is the density function of the lower of the two draws shoppers receive from F. Computing yields

$$\mathbb{E}_F\left[p_{(1,2)}\right] = \left(\frac{1-\alpha}{\alpha}\right)^2 \frac{v}{2}\left[\log\left(\frac{1-\alpha}{1+\alpha}\right) + \frac{2\alpha}{1-\alpha}\right].$$

Then $\mathbb{E}_F\left[p_{(1,2)}\right] < v\left(1-\alpha\right)$ can be rearranged to obtain

$$\log\left(\frac{1+\alpha}{1-\alpha}\right) > 2\alpha$$

which holds for $\alpha \in (0,1)$. ■

Corollary 1. *Within the mixed-price equilibrium firm responses of Lemma 1, as $c \in [0, v(1-\alpha)]$ increases, the expected price paid by both types of consumer increases.*

Proof: From Lemma 1 the pricing strategy for $c \in [0, v(1-\alpha)]$ is given by, $G(p;c)$ over $[\underline{p}(c), v]$. Differentiating,

$$\frac{dG\left(p;c\right)}{dc} = \frac{1}{c(n-1)-np}\left(\frac{\left(v-p\right)\left(1-\alpha\right)}{\alpha np - \alpha\left(n-1\right)c}\right)^{\frac{1}{n-1}}.$$

The second term is ≥ 0 else $G(p;c) > 1$. The first term is $\leq 0 \iff c\frac{n-1}{n} < p$ which is ensured because $p \geq c > c\frac{n-1}{n}$ when $c \leq v(1-\alpha)$ (this follows because $\underline{p}(c) \geq c \iff v(1-\alpha) \geq c$). Then for any $c, c' \in [0, v(1-\alpha)]$, if $c > c'$ then the equilibrium pricing distribution under c first order stochastic dominates that under c'.

Hence the expected price (paid by auto-renewers) and the expected lowest price from n draws (paid by shoppers) are higher under c. ■

Corollary 2. *As the number of firms increases, the expected price paid by shoppers decreases, but the expected price paid by auto-renewers increases.*

Proof: Using an observation from Morgan et al. (2006), industry profit of firms is given by $\alpha\mathbb{E}_G[p_{(1,n)}] + (1-\alpha)\mathbb{E}_G[p] = v(1-\alpha)$, where $\mathbb{E}_G[p_{(1,n)}]$ denotes the lowest price of n draws from $G(p; v(1-\alpha))$ and $\mathbb{E}_G[p]$ denotes the expected price from $G(p; v(1-\alpha))$. The RHS is the industry profit of firms if it charged v and only sold to auto-renewers. Differentiating and rearranging,

$$\frac{d\mathbb{E}_G[p_{(1,n)}]}{dn} = -\left(\frac{1-\alpha}{\alpha}\right)\frac{d\mathbb{E}_G[p]}{dn}$$

so the derivatives have opposite signs. Now show that $\frac{\mathbb{E}_G[p]}{dn} \geq 0$. To do this, show that $G(p, v(1-\alpha))$ is stochastically ordered in n:

$$\frac{dG(p; v(1-\alpha))}{dn} \leq 0$$

$$\iff \log\left(\frac{(v-p)(1-\alpha)}{\alpha(np - v(1-\alpha)(n-1)}\right) + \frac{(n-1)(p - v(1-\alpha))}{np - v(1-\alpha)(n-1)} \equiv X(p) \leq 0.$$

Note that $X(v(1-\alpha)) = 0$, and $\frac{dX}{dp} < 0$:

$$\frac{dX(p)}{dp} \leq 0 \iff p \geq v(1-\alpha)\left[\frac{(2n-2+\alpha(n-1)^2)}{(2n-1+\alpha(n-1)^2)}\right].$$

Notice that the term on RHS in square brackets is below 1. Recall that $p \geq v(1-\alpha)$ as this is the lower bound of the support, hence this is satisfied. ■

Proposition 4. *With $K > 1$ competing PCWs, if shoppers check $r = 1$ PCW then both types of consumer are worse off with the PCWs than without.*

Proof: There are no equilibria where $c \in [0, v(1-\alpha))$: Lemma A10 shows that there is a profitable upward deviation to $c_K \in (c_1, \underline{p}(c))$. It is profitable because there is an increase in fee-level but no reduction in the quantity of sales. Note that there can be no profitable deviation downwards given $r = 1$ because the fee would be reduced for no gain in the quantity of sales.

There are no equilibria where $c \in (v(1-\alpha), v]$: Lemma A9 shows $u_k = 0$ for all k. Lemmas A6 and A8 show that a deviation to $c_1 \in (0, v(1-\alpha))$ will generate $u_1 > 0$.

The only remaining option is $c_1 = \cdots = c_K \equiv c = v(1-\alpha)$, where firm responses are described by Lemma A11. If firms play a pure-pricing strategy, then as in the previous case, PCWs make zero profit and there is a profitable deviation to $c_1 \in (0, v(1-\alpha))$. If however, firms respond with the mixed-price strategy, this fee-level is an equilibrium: There can be no profitable deviation downwards for PCWs because the fee would be reduced for no gain in the quantity of sales as $r = 1$. Following an upward deviation from PCW$_K$ to $c_K \in (v(1-\alpha), v]$, when firms respond with a pure-pricing strategy as detailed in Lemma A11, PCW$_K$'s profit falls to zero. There is then an equilibrium at this fee level, and it is the unique such level where there exists an equilibrium.

At this fee-level, firms play just as they did when $K = 1$ with the same CDF over prices and all PCWs list all n firm prices. Both types of consumer are therefore left with the same level of surplus they had under $K = 1$. ■

Corollary 3. *When pricing is determined by $G(p; c)$, as the proportion of shoppers α increases, expected prices for both types of consumer decrease.*

Proof: Differentiating $G(p; c)$ by α,

$$\frac{dG\left(p;c\right)}{d\alpha} = \frac{1}{\alpha(n-1)(1-\alpha)}\left(\frac{\left(v-p\right)\left(1-\alpha\right)}{\alpha np - \alpha\left(n-1\right)c}\right)^{\frac{1}{n-1}}$$

The first term is > 0. The second term is always ≥ 0 or $G(p; c) > 1$. Then for any $\alpha, \alpha' \in (0, 1)$, if $\alpha > \alpha'$ then the equilibrium pricing distribution under α' first order stochastic dominates that under α. Hence the expected price (paid by auto-renewers) and the expected lowest price from n draws (paid by shoppers) are lower under α. ■

Corollary 4. *If the PCW can determine α as well as c in the preliminary stage, then the PCW sets $\alpha = \frac{1}{2}$.*

Proof: I expand the PCW's action set to include $\alpha \in (0, 1)$. Notice that for any choice of $\alpha \in (0, 1)$, by the reasoning as in the proof of Proposition 2, the PCW will avoid the pure equilibria of Lemma 1 so $c \in [0, v(1-\alpha)]$, firms mix and the PCW is given by

$c\alpha\left(\frac{n-1}{n}\right)$. The PCWs optimization problem can hence be solved by,

$$\max_{c,\alpha} \ c\alpha\left(\tfrac{n-1}{n}\right) \ \text{ s.t. } c \in [0, v(1-\alpha)] \text{ and } \alpha \in (0,1)$$

where the solution is $c = v(1-\alpha)$, $\alpha = \frac{1}{2}$. ■

A5. Results for K > 1 and r > 1

Lemma A13. *If $r > 1$, $c_1 = \cdots = c_{K-1} \in [0, v)$, $c_K \in (c_1, v]$, $u_K = 0$.*

Proof: This Lemma says that there is no profitable upwards deviation for a PCW from any equilibrium. Consider such a unilateral deviation by PCW$_K$. Firm response can either be pure or mixed pricing. If pure, then $u_k = 0$ for all k. If mixed, then at any price p that has positive probability of sales through PCWs where $K \in \mathcal{K}$, firms have a strict preference to play $(p, \{1, \ldots, K-1\})$ instead. This is because $r > 1$: Every shopper who sees the prices on PCW$_K$ also sees the prices on another PCW. Firms can therefore avoid PCW$_K$'s higher fee by not listing there while facing no reduction in the quantity of sales. ■

Lemma A14. *If $r = K$, $c_2 = \cdots = c_K \in (0, v(1-\alpha)]$ and $c_1 \in [0, c_2)$: $\beta_1 = 1$, $\beta_k = 0$ for $k = 2, \ldots K$ and prices are distributed according to the CDF $G(p, c_1)$.*

Proof: There are no point masses by Lemma A6. By Lemma A8 $\beta_1 = 1$. $\beta_k = 0$ for $k = 2, \ldots, K$ follows because all shoppers check every PCW ($r = K$). Therefore, at any price p that has positive probability of generating sales through PCWs, firms have a strict preference only to list on PCW$_1$ without any reduction in the quantity of sales. By Lemma A7 the maximum of the support is v. Firms must be indifferent to all $(p, 1)$ they are called upon to play, hence:

$$v\left(\tfrac{1-\alpha}{n}\right) = p\left(\tfrac{1-\alpha}{n}\right) + \tfrac{\alpha}{n}(1 - G(p; c_1))^{n-1}\left[np - (n-1)c_1\right].$$

which can be re-arranged to give $G(p, c_1)$ from (A2). ■

Proposition 5. *With $K > 1$ competing PCWs, shoppers are guaranteed to be better off than before the introduction of PCWs if and only if they check all PCWs: The unique equilibrium PCW fee-level is $c = 0$ if and only if $r = K > 1$.*

Proof: Sufficiency: Suppose not. Then there exists an equilibrium with $c > 0$. By Lemma A13, there is no profitable upward deviation. Now consider a downward deviation. If $c \in [v(1-a), v]$ and firms respond with a pure-pricing equilibrium, as detailed in Lemma A9, then $u_k = 0$. A deviation by PCW$_1$ to $c_1 \in (0, v(1-\alpha))$ would lead to the responses detailed in Lemma A14 and deviation profit of $u_1 = c\alpha\frac{n-1}{n} > 0$. If $c \in (0, v(1-a)]$ and firms respond with by mixing by $G(p,c)$, as in Lemma A14, then $u_k = c\frac{\alpha}{K}\frac{n-1}{n} > 0$ for all k. But a deviation by PCW$_1$ to $c_1 < c$ exists s.t. $u_1 = c_1\alpha\frac{n-1}{n} > c\frac{\alpha}{K}\frac{n-1}{n} = u_k$. This deviation is strictly profitable, so c could not have been an equilibrium. To see that $c = 0$ is an equilibrium, recall that by Lemma A13 there is no profitable upward deviation.

Necessity: Lemma A11 shows that for $K \geq r = 1$ the unique equilibrium fee level is $c = v(1-\alpha)$. Lemma 2 shows that for $1 < r < K$ there are multiple equilibria. Hence $K = r > 1$ is the only case where the unique equilibrium of $c = 0$ obtains.

Consumer welfare: As noted in the text, as firms make the same expected profit in both worlds, there is a one-to-one relation between consumer welfare and PCW profit. To see the difference in the changes to shopper and auto-renewer welfare from a move to a world with a PCW but $c = 0$, Proposition 3 of Morgan et al. (2006) (the only difference is that they have $v = 1$) shows that the increase from q to n results in a reduction in the expected price paid by shoppers, and an increase for auto-renewers. ■

Lemma 2. *When $K > r > 1$, there exists an equilibrium in which PCWs charge $\bar{c} > 0$, firms list on all PCWs, mixing over prices by $G(p; \bar{c})$ where*

$$\bar{c} = \frac{v(1-\alpha)Kr(K-r)}{K(1+r(K-2)) + \alpha r(K-1)(r-1)(n-1)}.$$

Proof: I show that there exist equilibria such that $c \in [0, \bar{c}]$. Note that $\bar{c} \in [0, v(1-\alpha))$ hence any $c \in [0, v(1-\alpha))$. Take such a c as a candidate equilibrium fee level. There are no point masses by Lemma A6. By Lemma A7 the maximum of the support is v. By Lemma A5 $pr(E) = 0$ and as all PCWs charge the same fee, firms are content to list in all of them i.e., $\beta_k = 1$ for all k. Firms must be indifferent to all $(p, \{1, \ldots, K\})$

they are called upon to play, hence

$$v\left(\tfrac{1-\alpha}{n}\right) = p\left(\tfrac{1-\alpha}{n}\right) + \tfrac{\alpha}{n}(1-G(p;c))^{n-1}\left[np-(n-1)c\right]$$

which can be re-arranged to give $G(p;c)$ in (A2).

To confirm c is an equilibrium fee level, ensure there is no profitable PCW deviation. By Lemma A13, there is no profitable upward deviation. However, unlike Lemma A14, when $1 < r < K$ it is no longer true that any undercut by PCW_1 to $c_1 < c$ will result in all firms listing on it exclusively. This is because consumers see $r > 1$ PCWs: when a firm sells having played $(p, \{1, \ldots, K\})$ it pays $c_1\frac{1}{K} + c\frac{K-1}{K}$, but were it to have played $(p, 1)$ it would have paid $c_1\frac{r}{K}$ i.e., the firm faces a trade-off between lower fees and higher sales volume. By Lemma A8 $\beta_1 = 1$. PCW_1's deviation profit is therefore determined by β_k, $k > 1$. For small-enough undercuts of c, firms will still be content to list on all PCWs. Suppose however, that PCW_1 undercuts by just enough such that such that firms are no longer content to list all their prices on the other PCWs. In the best case for PCW_1, $\beta_k = 0$ for all k so that all shoppers checking PCW_1 buy only through PCW_1. If in this case, PCW_1 still does not make more than its equilibrium profit $u^* = c\frac{\alpha}{K}\frac{n-1}{n}$, then the undercut is never profitable and c constitutes an equilibrium fee level. I now carry out this logic.

Denote $c_1 < c$ as the undercut of PCW_1 and $\tilde{c}_1$ as the threshold level required for c_1 to entice firms to de-list some of their prices from PCWk s.t. $k > 1$. Given $\mathbf{c} = (c_1, c, \ldots, c)$, similarly to the derivation above, one can show firms will respond by $G(p;\mathbf{c})$ over $[\underline{p}(\mathbf{c}), v]$ as in (A2). Firm deviation profit from this response to $(p, 1)$ is given by

$$\pi(p) = p\left(\frac{1-\alpha}{n}\right) + \frac{(v-p)(1-\alpha)\left[Kp - r(p-c_1)(n-1)\right]}{n\left(Knp-(n-1)(c_1+c(K-1))\right)}.$$

Note that this is valid for $p \in [\underline{p}(\mathbf{c}), v]$, but $p < \underline{p}_1$ give strictly less profit than $p = \underline{p}(\mathbf{c})$. One can show that $\pi(p)$ is convex in p. Together with the observation that $(v, 1)$ gives the equilibrium profit $\frac{v(1-\alpha)}{n}$, this says that the optimal deviation is to $(\underline{p}(\mathbf{c}), 1)$ and is profitable if and only if $\pi(\underline{p}(\mathbf{c})) > \frac{v(1-\alpha)}{n}$, which can be rearranged to give

$$c_1 < \frac{c(K-1)(K+r\alpha(n-1)) - Kv(K-r)(1-\alpha)}{r\alpha(K-1)(n-1)+K(r-1)} \equiv \tilde{c}_1.$$

Suppose that the firm response following the undercut was such that $\beta_k = 0$ for all $k > 1$ when PCW_1 sets the highest such undercut just below $\tilde{c}_1$. PCW_1 prefers not to make the deviation whenever

$$u^* \geq \tilde{c}_1 \frac{\alpha r}{K}\frac{n-1}{n} \iff \frac{c}{r} \geq \tilde{c}_1$$

which can be rearranged to give

$$c \leq \frac{v(1-\alpha)Kr(K-r)}{K(1+r(K-2)) + \alpha r(K-1)(r-1)(n-1)} \equiv \bar{c}.$$

Therefore at fee levels $c \in [0, \bar{c}]$ there exist equilibria where $\beta_k = 1$ for all k and firms mix by the CDF given in the Lemma. ■

Proposition 6. *For any r, there exists a $\tilde{K}$ such that as long as there are more than $\tilde{K}$ aggregators both types of consumer can be worse off than before the introduction of any PCWs.*

Proof: To show for shoppers, let the equilibrium be given as in Lemma 2 with $c = \bar{c}$ and see that

$$\lim_{K\to\infty} \bar{c} = v(1-\alpha). \tag{A4}$$

For shopper welfare notice that for $K = r$, $\bar{c} = 0$ and

$$\mathbb{E}_F[p_{(1,q)}] \geq \mathbb{E}_{G(0)}[p_{(1,n)}] \tag{A5}$$

where $G(c)$ and F denote $G(p; c)$ and $F(p)$. This is because the PCWs effectively increase the number of firms competing for each shoppers from q to n, and from the last point of the proof of Proposition 5, this lowers the price paid by shoppers. Proposition 3 shows that both types are worse off with PCWs when $c = v(1-\alpha)$ for $K \geq r = 1$, but notice that if $c = v(1-\alpha)$ in equilibrium with $K > r \geq 1$, then this will be true a fortiori, because it can be shown that $G(v(1-\alpha))$ first-order stochastically dominates F for all q. By (A4) then as $K \to \infty$, $\mathbb{E}_F[p_{(1,q)}] < \mathbb{E}_{G(v(1-\alpha))}[p_{(1,n)}]$. Because $\bar{c}$ is continuous in K, $\frac{d\bar{c}}{dK} > 0$ and $\frac{d\mathbb{E}_{G(\bar{c})}[p_{(1,n)}]}{dK} > 0$ by Corollary 1, there exists $\tilde{K}$ s.t. for $K > \tilde{K}$ there is always an equilibrium fee level that makes shoppers worse off relative to a world without PCWs.

The proof for auto-renewers is similar but simpler to the above steps for shoppers and has p replacing $p_{(1,q)}$ and $p_{(1,n)}$: From Proposition 3 of Morgan et al. (2006), the effect of increasing the number of firms competing for each shopper increases prices for auto-renewers. Then by Corollary 1, their prices only rise further under $G(c)$. Hence auto-renewers are worse off for any c and hence any K. ■

A6. Results under Publicly Observable Fees

Definition (Coordinated Subgame)**.** *Given* **c***, a 'coordinated subgame' is when shoppers attend and firms list in only the cheapest PCWs i.e., shoppers in PCW$_k$ if $c_k = c_1$, and for firms $\beta_k = 1$ if $c_k = c_1$ else $\beta_k = 0$. Firms mix by $G(p, c_1)$.*

Lemma A15. *When fees are observed by shoppers, and all subgames are 'coordinated subgames' there exists a subgame perfect equilibrium where $c = 0$, $\beta_k = 1$ for all k and prices are distributed according to the CDF $G(p, 0)$.*

Proof: There are no profitable deviations for firms: $G(p, 0)$ ensures they are indifferent between listing all prices in the support $[\underline{p}(0), v]$, listing any price less than $\underline{p}(0)$ is less profitable than listing v, and any $(p, \emptyset)$ generates at most $\frac{v(1-\alpha)}{n}$ when $p = v$. To see there is no profitable upward deviation for PCWs, note that due to the 'coordinated subgame' shoppers and firms would not attend this PCW following such a deviation, so the PCW receives zero profit. To see that these coordination subgames are Nash equilibria: for firms one can conduct the same round of checks as at the beginning of this proof; for consumers notice that as the non-cheapest PCWs are all empty, there is no incentive to check them. ■

Proposition 7. *When $K > 1$ and PCW fees are observed by shoppers, there exists an equilibrium with $c = 0$.*

Proof: That there exists an equilibrium with $c = 0$ follows directly from Lemma A15. For shopper welfare at $c = 0$, see the last point of the proof of Proposition 5. ■

A7. Results under Price Discrimination

Now assume that different prices can be charged directly and through the PCW. Denote such prices p_0 and p_k respectively, where $k = 1, \ldots, K$ indexes the PCW as before.

Lemma A16. *If $r = 1$, $c_1 = \cdots = c_{K-1} = v(1-\alpha)$ and $c_K \in [c_1, v]$ then there exist firm responses s.t. $p_0 = p_1 = \cdots = p_K = v(1-\alpha)$ and $\beta_k = 1$ for all k.*

Proof: Under these strategies, firm profit is $\pi = \frac{v(1-\alpha)}{n}$. There is no profitable deviation involving p_k for any k: lower would prompt sales at $p_k < c_k$, higher would never attract any shoppers. There is no profitable deviation involving p_0: lower would still sell to auto-renewers and own-shoppers but at a lower price, higher would only sell to auto-renewers for which the optimal such deviation is to v generating profit $\frac{v(1-\alpha)}{n} = \pi$. ■

Proposition 8. *With price discrimination, if $r = 1$, there exists an equilibrium in which PCWs set $c = v(1-\alpha)$, firms list on all PCWs, $p_0 = v$ and $p_1 = \cdots = p_K = v(1-\alpha)$.*

Proof: Given $c_1 = \cdots = c_K \equiv c = v(1-\alpha)$, let us confirm there are no profitable deviations for firms. Equilibrium firm profit is $\pi = \frac{v(1-\alpha)}{n}$. There is no profitable deviation involving p_k for any k: lower would prompt sales at $p_k < c_k$, higher would never attract any shoppers. There is no profitable deviation involving p_0: lower would either sell only to auto-renewers at a lower price, or to both auto-renewers and own-shoppers but only for prices at least as low as $v(1-\alpha)$, generating profit no greater than $\frac{v(1-\alpha)}{n}$; higher would be above v and hence make zero profits. Now consider a PCW deviation. Equilibrium PCW profit is $u_k = \frac{\alpha v(1-\alpha)}{K} > 0$. There is no profitable deviation to a lower fee: as $r = 1$ doing so would at best sell to the same proportion of shoppers ($\frac{\alpha}{K}$) but at a lower fee. As for a deviation to a higher fee: assume that firms' mutual best responses are given by those in Lemma A16 where PCW profit is zero. ■

Corollary 5. *When firms price discriminate, following $c_1 \in (0, c_2)$ and $c_2 = c_3 = , \ldots, = c_K \equiv c \in (0, v(1-\alpha)]$ there exist mutual best-responses of firms such that they list on all PCWs, setting $p_0 = v$ and $p_k = \cdots = c_k$ for all k.*

Proof: Given $\mathbf{c}$ as in the Lemma, let us confirm there are no profitable deviations for firms. Firm profit is $\pi = \frac{v(1-\alpha)}{n}$. There is no profitable deviation involving p_k for any k: lower could prompt sales, but only at $p_k < c_k$, higher would never attract any shoppers. There is no profitable deviation involving p_0: lower would either sell only to auto-renewers at a lower price, or to both auto-renewers and own-shoppers but only for prices at least as low as c_2, generating profit no greater than $\frac{c_2}{n} < \pi$; higher would be above v and hence make zero profits. ■

Lemma A17. *If $r > 1$, $c_1 = \cdots = c_{K-1} = 0$ and $c_K \in [0, v]$ then there exist firm responses s.t. $p_0 = v$, $p_1 = \cdots = p_{K-1} = 0$, $\beta_k = 1$ for $k < K$ and $\beta_K = 0$.*

Proof: Under these strategies, firm profit is $\pi = \frac{v(1-\alpha)}{n}$. There is no profitable deviation involving p_K: $p_K < 0$ would give sell at a loss whereas $p_K > 0$ would never attract any shoppers as $r > 1$. There is no profitable deviation involving p_k for any $k < K$: Lower would only create sales at a loss, higher would never attract any shoppers. There is no profitable deviation involving p_0: lower would still sell to auto-renewers and (for $p_0 \leq 0$) own-shoppers but at a lower price, higher would sell to no-one. ■

Lemma A18. *If $r > 1$, there exists an equilibrium where $c = 0$, firms list on all PCWs, $p_0 = v$ and $p_1 = \cdots = p_K = 0$..*

Proof: One can follow the proof of Proposition 8 to confirm there are no profitable deviations for firms. For PCWs, one only need consider an upward deviation in fee level in which case, assume firms respond as in Lemma A17, which yields zero profit for the deviating PCW. ■

Proposition 9. *For $n \to \infty$, following the introduction of a PCW, the expected prices faced by both types of consumer in the setting without price discrimination (Proposition 2) are the same as in the setting with price discrimination (Proposition 8).*

Proof: Taking limits,

$$\lim_{n\to\infty} G(p; c) = 0$$

which shows $\lim\limits_{n\to\infty} \mathbb{E}_G[p] = v$. The CDF for the lowest of n draws is given by

$$H(p; c) = 1 - (1 - G(p; c))^n = 1 - \left[\frac{(v-p)(1-\alpha)}{\alpha\,(np - (n-1)c)}\right]^{\frac{n}{n-1}}$$

for $p \in [\underline{p}(c), v]$. Taking limits,

$$\lim_{n\to\infty} H(p; c) = 1 \text{ and } \lim_{n\to\infty} \underline{p}(c) = c$$

which shows $\lim\limits_{n\to\infty} \mathbb{E}_G[p_{(1,n)}] = c$. These prices are the same as those in the equilibrium of Proposition 8. ■

A8. Results with Search Costs

In the world without a PCW under $q = 2$, one can compute

$$\begin{aligned}
\mathbb{E}_F[p] &= \frac{v}{2}\left[\frac{1-\alpha}{\alpha}\right]\left[\log\left(\frac{1+\alpha}{1-\alpha}\right)\right] \\
\mathbb{E}_F[p_{(1,2)}] &= \frac{v}{2}\left[\frac{1-\alpha}{\alpha}\right]^2\left[\log\left(\frac{1-\alpha}{1+\alpha}\right) + \frac{2\alpha}{1-\alpha}\right] \\
\mathcal{B}_0 &= \frac{v}{2}\left[\frac{1-\alpha}{\alpha^2}\right]\left[\log\left(\frac{1+\alpha}{1-\alpha}\right) - 2\alpha\right].
\end{aligned}$$

In the world with a PCW under $n \to \infty$, Proposition 9 shows that,

$$\mathbb{E}_G[p] = v, \qquad \mathbb{E}_G[p_{(1,\infty)}] = c, \qquad \mathcal{B}_1 = v - c$$

where $c = v(1 - \tilde{\alpha}) \leq v(1 - \alpha)$ in equilibrium.

Proposition 10. *Shoppers, converts and non-converts can all be worse off with a PCW than without.*

Proof: This is a proof by example. Let $v = 1$, $\alpha = 0.7$, $n = 3$ and $q = 2$. Then, in the world without a PCW, it can be computed that,

$$\mathbb{E}_F[p] = 0.3717, \qquad \mathbb{E}_F[p_{(1,2)}] = 0.2693, \qquad \mathcal{B}_0 = 0.1024.$$

In the world with a PCW, assume $\tilde{\alpha} = 0.71$ (i.e., the PCW attracts 0.01 more consumers i.e., $\frac{1}{30}$ auto-renewers, to compare prices), so $c = v(1 - \alpha) = 0.29$ and

$$\mathbb{E}_G[p] = 0.5557, \qquad \mathbb{E}_G[p_{(1,3)}] = 0.3748, \qquad \mathcal{B}_1 = 0.1808.$$

Comparing, one can see that if there is an S such that this can be rationalized as an equilibrium, then all types of consumer will be worse off with a PCW than without. To rationalize, I construct an S such that:

$$S(\mathcal{B}_1) = \begin{cases} 0 & \text{if } \mathcal{B}_1 < \underline{s} \\ \dfrac{\mathcal{B}_1 - \underline{s}}{\mathcal{B}_1} & \text{else} \end{cases}$$

where $\underline{s}$ is determined by

$$S(0.1808) = \frac{1}{30} \iff \underline{s} = 0.1748.$$

Finally, notice that $\underline{s} > \mathcal{B}_0$. ■

Lemma A19. *Under maximum potential, any $\tilde{\alpha}$ can be rationalized by an S.*

Proof: Note that any $\tilde{\alpha}$ can be rationalized with an S (as in the proof of Proposition 10) because $\mathcal{B}_0 < \mathcal{B}_1$. To see this holds, one can show that $\mathcal{B}_0 < v\alpha$ holds for any $\alpha \in (0, 1)$ and hence so does $\mathcal{B}_0 < v\tilde{\alpha}$. ■

Proposition 11. *If the market has maximum potential: converts are better off; shoppers can be worse off; and non-converts are worse off with a PCW than without.*

Proof: As $v(1 - \alpha) > v(1 - \tilde{\alpha})$ when there are converts, to show converts are better off with a PCW under maximum potential one can show,

$$\mathbb{E}_F[p] > v(1 - \alpha) \iff \log\left(\frac{1 + \alpha}{1 - \alpha}\right) > 2\alpha$$

which is satisfied for $\alpha \in (0, 1)$. Note that $\tilde{\alpha}$ is rationalizable by some S (Lemma A19).

Shoppers are worse off under maximum potential wherever $v(1 - \tilde{\alpha}) > \mathbb{E}_F[p_{(1,2)}]$. To show this can occur, note that $v(1 - \alpha) > \mathbb{E}_F[p_{(1,2)}]$ holds for all $\alpha \in (0, 1)$ (see Proposition 3). Hence for any α there exists some $\tilde{\alpha} > \alpha$ small enough such that shoppers are worse off. Note that such a $\tilde{\alpha}$ is rationalizable by some S (Lemma A19).

Non-converts are worse off under maximum potential because $\mathbb{E}_F[p] < \mathbb{E}_G[p] = v$.

■

References

Admiral Group PLC (2012), “Annual Report.”

Admiral Group PLC (2013), “Annual Report.”

Armstrong, Mark (2006), “Competition in two-sided markets.” *RAND Journal of Economics*, 37, 668–691.

Armstrong, Mark and Jidong Zhou (2011), “Paying for prominence.” *The Economic Journal*, 121, F368–F395.

Arnold, Michael A, Chenguang Li, Christine Saliba, and Lan Zhang (2011), “Asymmetric market shares, advertising and pricing: Equilibrium with an information gatekeeper.” *Journal of Industrial Economics*, 59, 63–84.

Arnold, Michael A and Lan Zhang (2014), “The unique equilibrium in a model of sales with costly advertising.” *Economics Letters*, 124, 457–460.

Baye, Michael R, Xiaxun Gao, and John Morgan (2011), “On the optimality of click-through fees in online markets.” *Economic Journal*, 121, F340–F367.

Baye, Michael R, Dan Kovenock, and Casper G de Vries (1992), “It takes two to tango: Equilibria in a model of sales.” *Games and Economic Behavior*, 4, 493–510.

Baye, Michael R and John Morgan (2001), “Information gatekeepers on the internet and the competitiveness of homogeneous product markets.” *American Economic Review*, 91, 454–474.

Baye, Michael R and John Morgan (2009), “Brand and price advertising in online markets.” *Management Science*, 55, 1139–1151.

Baye, Michael R, John Morgan, and Patrick Scholten (2004), “Price dispersion in the small and in the large: Evidence from an internet price comparison site.” *Journal of Industrial Economics*, 52, 463–496.

Baye, Michael R, John Morgan, and Patrick Scholten (2006), “Information, search, and price dispersion.” In *Handbook on Economics and Information Systems* (T Hendershott, ed.), volume 1, Elsevier, Amsterdam.

Belleflamme, Paul and Martin Peitz (2010), “Platform competition and seller investment incentives.” *European Economic Review*, 54, 1059–1076.

BGL Group Ltd. (2013), “Group Results and Review.”

Brown, Jeffrey R. and Austan Goolsbee (2002), “Does the internet make markets more competitive? Evidence from the life insurance industry.” *Journal of Political Economy*, 110, 481–507.

Brynjolfsson, Erik and Michael D Smith (2000), “Frictionless commerce? A comparison of internet and conventional retailers.” *Management Science*, 46, 563–585.

Burdett, Kenneth and Kenneth L Judd (1983), “Equilibrium price dispersion.” *Econometrica*, 51, 955–969.

Byrne, David P., Nicolas de Roos, and Daniel Tiong (2014), “The internet, search and asymmetric pricing: A natural experiment in retail gasoline.” Working Paper, University of Melbourne.

Caillaud, Bernard and Bruno Jullien (2003), “Chicken & egg: Competition among intermediation service providers.” *RAND journal of Economics*, 34, 309–328.

Chioveanu, Ioana (2008), “Advertising, brand loyalty and pricing.” *Games and Economic Behavior*, 64, 68–80.

Consumer Focus (2013), “Comparing comparison sites.”

de Cornière, Alexandre and Greg Taylor (2014), “Quality provision in the presence of a biased intermediary.” Working Paper.

Diamond, Peter A (1971), “A model of price adjustment.” *Journal of Economic Theory*, 3, 156–168.

Edelman, Benjamin G and Julian Wright (2015), “Price coherence and excessive intermediation.” *Quarterly Journal of Economics*, 130, 1283–1328.

Ellison, Glenn and Sara Fisher Ellison (2009), "Search, obfuscation, and price elasticities on the internet." *Econometrica*, 77, 427–452.

Ellison, Glenn and Drew Fudenberg (2003), "Knife-edge or plateau: When do market models tip?" *Quarterly Journal of Economics*, 118, 1249–1278.

Ellison, Glenn and Alexander Wolitzky (2012), "A search cost model of obfuscation." *RAND Journal of Economics*, 43, 417–441.

Fraser, Douglas (2014), "Heat is on: Why are energy bills so high?" *BBC News*, URL http://www.bbc.com/news/uk-scotland-scotland-business-25743336.

Goodley, Simon and Gwyn Topham (2014), "Go compare founder to bank £44m after selling remaining stake to esure." *The Guardian*, URL http://www.theguardian.com/business/2014/dec/08/go-compare-founder-bank-44m-selling-remaining-stake-esure.

Gorodnichenko, Yuriy, Viacheslav Sheremirov, and Oleksandr Talavera (2015), "Price setting in online markets: Does IT click?" University of California, Berkeley Working Paper.

Hagiu, Andrei (2009), "Two-sided platforms: Product variety and pricing structures." *Journal of Economics & Management Strategy*, 18, 1011–1043.

Kessler, Aaron M and Julie Weed (2015), "Hotels fight back against sites like Expedia and Priceline." *New York Times*, URL http://www.nytimes.com/2015/09/01/business/hotels-direct-booking-online-travel-sites.html?_r=0.

Milligan, Brian (2015), "Energy comparison sites can earn £30 per switch, MPs told." *BBC News*, URL http://www.bbc.com/news/business-31110625.

Mintel Group Ltd. (2013), "Web aggregators in financial services."

Moneysupermarket Group PLC (2013), "Annual Report."

Moraga-González, José L, Zsolt Sándor, and Matthijs R Wildenbeest (2015), "Prices and heterogeneous search costs." Working Paper, University of Amsterdam.

Morgan, John, Henrik Orzen, and Martin Sefton (2006), "An experimental study of price dispersion." *Games and Economic Behavior*, 54, 134–158.

Poulter, Sean (2015), "Hotel guests 'are being fleeced by online agents': Hoteliers say they are forced to hike prices to cover 25% commission taken by websites." *The Daily Mail*, URL http://www.dailymail.co.uk/news/article-3044298/Hotel-guests-fleeced-online-agents-Hoteliers-say-forced-hike-prices-cover-25-commission-taken-websites.html.

Reisinger, Markus (2014), "Two-part tariff competition between two-sided platforms." *European Economic Review*, 68, 168–180.

Rickard-Straus, Rachel (2015), "We can't act surprised that comparison sites make a mint from us switching - Snoop Dogg doesn't work for free." *This is Money*, URL http://www.thisismoney.co.uk/money/bills/article-2933401/Energy-price-comparison-sites-spend-110m-annoying-adverts.html.

Rochet, Jean-Charles and Jean Tirole (2003), "Platform competition in two-sided markets." *Journal of the European Economic Association*, 1, 990–1029.

Rosenthal, Robert W (1980), "A model in which an increase in the number of sellers leads to a higher price." *Econometrica*, 48, 1575–1579.

Salop, Steven and Joseph Stiglitz (1977), "Bargains and ripoffs: A model of monopolistically competitive price dispersion." *Review of Economic Studies*, 44, 493–510.

Shearman, Sarah (2011), "Adspend is key to price comparison sites' success." URL http://www.marketingmagazine.co.uk/article/1058209/adspend-key-price-comparison-sites-success.

Stahl, Dale O (1989), "Oligopolistic pricing with sequential consumer search." *American Economic Review*, 79, 700–712.

Stahl, Dale O (1996), "Oligopolistic pricing with heterogeneous consumer search." *International Journal of Industrial Organization*, 14, 243–268.

Stigler, George J (1961), "The economics of information." *Journal of Political Economy*, 69, 213–225.

Varian, Hal R (1980), "A model of sales." *American Economic Review*, 70, 651–659.

Multi-Attribute Decision by Sampling: An Account of the Attraction, Compromise and Similarity Effects

1. Introduction

How do we choose between a small apartment close to our workplace and a larger apartment with a longer commuting distance? How do we weigh the price of a hotel against its Hotels.com quality rating, and how do we select between two different employees when personality favors one but sales figures favor the other? Such choices, involving objects that differ on multiple attributes, are ubiquitous in everyday life.

A particular challenge to models of multi-attribute choice is given by three much-studied context effects, each of which is of both theoretical and practical importance: the similarity effect (Tversky, 1972), the attraction effect (Huber *et al.*, 1982) and the compromise effect (Simonson, 1989). Initial explanations focused on one or two of these effects at a time, making reference to decision strategies such as elimination-by-aspects (Tversky, 1972) or concepts such as loss aversion (Simonson and Tversky, 1992). Models have however typically had difficulty in accounting for all three effects within a unifying framework without resorting to arguably ad hoc parameters or separate mechanisms in order to capture all three effects simultaneously. Here we offer a concise account based on a single mechanism: Multi-Attribute Decision by Sampling (MADS). MADS combines and extends three independently-developed ideas: (i) that individuals compare targets with comparison items; (ii) that comparisons are based on simple dominance relations; and (iii) that comparison items are systematically influenced by the choice set faced.

According to MADS, the evaluation of an alternative is based on how it fares *relative* to comparators in the mental choice context. Faced with a choice set, an individual draws a sample of items to mind. A choice alternative then accrues a point for every item in the sample (and every other item in the choice set) that it dominates.[1] The

[1] If many choice alternatives dominate a sampled item, one of the dominating alternatives is selected at random to accrue a point.

alternative with the highest score is then selected.

We test the model's explanation of context effects with two experiments. In the first, we manipulate choice. The treatment exposes individuals to a selection of items that are dominated by some of the choice alternatives in order to affect the items that are brought to mind during choice. The overall effect was that alternatives promoted by our treatment were chosen 7% more often. Our theory is based on simple dominance relations where an item dominates another when it is considered better on all its attributes. We use two-attribute choice items (hotels that vary in price and quality) and find that the treatment effect can be decomposed as 17% for those participants who reported paying equal attention to both attributes (the majority), whereas there was no effect for those reporting unequal attention.

The choice manipulations fed through to impact the size and significance of the classic context effects. Our design allowed us to attempt to enhance and counter the effects. Doing so, our treatment is shown to cause the attraction effect to be doubled and then to be reduced by enough to be insignificantly different from zero i.e., the attraction effect is turned off. The compromise and similarity effect were also significantly manipulated, but by less than the attraction effect. In particular, although the similarity effect was not replicated, it was enhanced by enough to be significantly difference from zero i.e., the similarity effect was turned on.

In a second experiment, we elicited the distributions which MADS assumes individuals sample comparison items from. We show that these distributions (aggregated over individuals) depend on the choice faced. Furthermore, the dependence of these aggregate distributions is such that the model predicts the context effects to occur in expectation. We then examine individual-level distributions and show that MADS predicts the choices in most of our conditions. To predict individual choice, we also estimate the model's central parameter: the number of comparison items brought to mind. We find this to be four, which we note is a psychologically realistic value for the capacity of human working memory (Cowan, 2001).

The paper is organized as follows. After a review of the attraction, similarity, and compromise effects, we note that extant cognitive and economic models are limited in their ability to provide a unified account in which all three effects emerge naturally without the need for additional effect-specific mechanisms. We then provide an informal intuitive account of how MADS accounts for context effects. We then report two experiments to test the predictions of MADS that (a) changing the distribution of background stimuli (hotels varying in price and quality rating) seen by participants will change their choices in ways predicted by MADS, thus enabling us to counter or enhance all three context effects in predictable ways, and (b) participants' inferred distributions will be changed in predictable ways by the context of choice options confronting participants.

2. Context Effects

We illustrate the three context effects in Figure 1 where items are located within two-dimensional price$\times$quality space. Consider the low-quality, low-price option A, and the higher-quality, higher-price option B. An attraction effect occurs when one of two options is more likely to be chosen when a third option that is dominated by it (but not dominated by the other) is introduced, e.g., $p(A|\{A, B, T_A\}) > p(A|\{A, B, T_B\})$. A compromise effect occurs when an option is more likely to be chosen when it becomes an intermediate option, e.g., $p(B|\{A, B, C_B\}) > p(B|\{A, B, C_A\})$. A similarity effect occurs when a third option that is similar to one of the alternatives increases the probability of choosing the dissimilar alternative, e.g., $p(A|\{A, B, S_A\}) > p(A|\{A, B, S_B\})$.

Figure 1: Context effects

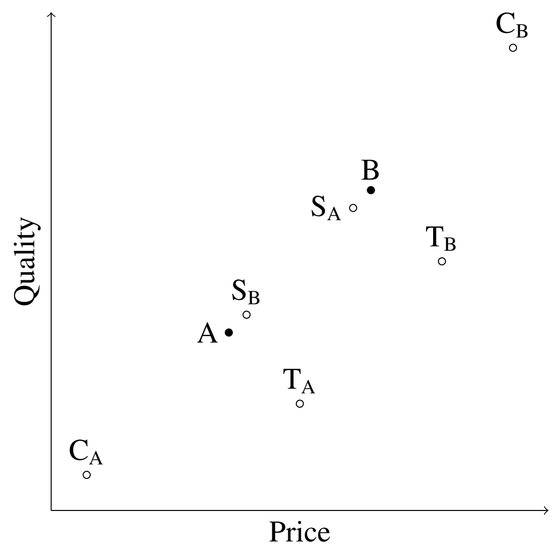

Note here that because a higher price is less preferred, indifference curves have positive slopes. Hollow dots represent the various decoys that join A and B to make up ternary context-effect choice sets.

These "big three" context effects are found when psychophysical stimuli are used (Trueblood *et al.*, 2013; Tsetsos *et al.*, 2012) in memory retrieval (Maylor and Roberts, 2007), inference tasks (Trueblood, 2012), and in-store purchasing (Doyle *et al.*, 1999) as well as the lab (e.g., Pettibone and Wedell, 2000; Wedell, 1991).

3. Literature

Several models of the three context effects have been put forward in both psychology and economics. For example, attraction and compromise effects are attributed to loss aversion in the Leaking Competing Accumulators model (Usher and McClelland, 2004) or attention switching and mutual inhibition occurring between choice options in Multialternative Decision Field Theory (Roe *et al.*, 2001). Bhatia (2013) proposes a model in which the accessibility of attributes is determined by the attributes' associations with objects of potential choice. More accessible attributes in turn carry higher weight in an evidence accumulation process. These models, among others, can all account for the three key context effects. However, in all cases the similarity effect does not arise for the same reason as the attraction and compromise effects. In Bhatia's model for example, the three effects can occur simultaneously but will not do so under all parameter settings. In the Multiattribute Linear Ballistic Accumulator model (Trueblood *et al.*, 2014), the attraction and compromise effects arise because objects that are closer to each other receive larger attention weightings, whereas the similarity effect occurs when confirmatory evidence is given more weight than disconfirmatory evidence. Models developed within economics also do not offer an account of all three effects (e.g., Kamenica, 2008).

Here we offer an alternative approach, aiming to provide a unifying account in which all three context effects emerge from the same mechanism without the need for separate parameters. Our model has one free parameter to be estimated: the number of comparator items brought to mind.

Our approach is rooted in and brings together various approaches within cognitive science, consumer psychology and economics. In economics, some recent theoretical approaches have been developed to show how choice behavior can be explained by cognitive limitations such as binary ordinal comparison (e.g., Kornienko, 2013), memory limitations (Mullainathan, 2002), psychological salience (Bordalo *et al.*, 2012, 2013) or as optimal responses to noise (e.g., Robson, 2001; Steiner and Stewart, 2014). Our approach falls within this tradition and also draws on information-sampling models of judgment, most of which assume that judgments are typically made on the basis of limited samples (e.g., Fiedler, 2000; Fiedler and Juslin, 2006; Fiedler and Kareev, 2006; Lindskog *et al.*, 2013). Relevant alternatives are assumed typically to be retrieved from memory as well as, or instead of, being sampled from the choice context (for a related approach within economics, see Gennaioli and Shleifer, 2010). A direct precursor of our work within psychology is the Decision by Sampling model (DbS: Stewart *et al.*, 2006, 2015); we return to this below.

Finally, recent work within psychology emphasizes that estimates of relevant background distributions and samples may themselves be influenced by choice options (e.g., Sher and McKenzie, 2014). Within economics, recent models accommodate the fact that available choice options convey relevant information about the market (e.g., Kamenica, 2008). Within consumer psychology, it has also been suggested that people take available choice options as informative about the marketplace. For example, a large person will likely (and wisely) choose a sweatshirt size near the top of the available range of size options (Prelec *et al.*, 1997; Simonson, 2008; Wernerfelt, 1995). Primes that change individuals' estimate of their own ranked position within the population may similarly change their choice of product: Individuals given easier golf putting tasks subsequently chose higher-level golfing equipment (Gershoff and Burson, 2011). More generally, the idea that people update their estimates about attribute values on the basis of experimentally-provided options has received traction in cognitive psychology. This idea has been used to explain preference reversals (Sher and McKenzie, 2014), context effects of the type discussed in the present paper (Shenoy and Yu (2013)), effects of distribution skew on judgments (Brown *et al.*, 2015b) and anchoring effects (Brown *et al.*, 2015a). Next, we describe how MADS builds on these various approaches.

4. Multi-Attribute Decision by Sampling

Before defining the model formally, we first provide an intuitive introduction to the model. Following the formal definition, we give a simple graphical illustration of how the model accounts for the attraction, similarity and compromise effects.

In Multi-Attribute Decision by Sampling, the evaluation of an alternative x is based on how it fares *relative* to comparators in the mental choice context. More precisely, we consider x's value to be determined by its position within the cumulative distribution of alternatives inferred following exposure to a choice set. By definition, the cumulative distribution gives the proportion of alternatives which x dominates, ranging between 0 and 1. This gives a simple measure of the relative position of x. The conceptualization of the value of an item being determined by its location in the market distribution as measured by its CDF position is intimately related to the fundamental notion that consumers prefer items more if they considered more of a bargain. Indeed if the degree to which an item is considered a bargain is defined as how many other items in the market it is unambiguously better than, then this is exactly our notion of value.

In process terms, we assume that individuals bring to mind $n > 0$ comparators through recall or inference about the product space. Individuals then use these comparators to estimate an alternative's CDF value by counting the number of comparison items it dominates. We assume the alternative with the highest such value is chosen. Note that because the set of comparison items is finite, whether we use counts or proportions for value is irrelevant for selecting the highest-value item. We use counts throughout the paper, but note that as $n \to \infty$, it is the proportion (i.e., count divided by n) which approaches x's position in the inferred cumulative distribution.

To illustrate the choice procedure consider an alternative x in choice set X. It accumulates a score (or count) $s(x)$ based on how many alternatives in the comparison sample (and other alternatives in the choice set) it dominates. For example, consider hotel stays differing in price and quality. If the choice set is $\{A, B\} = \{(\$100, 4.0), (\$200, 4.5)\}$ and one other hotel is brought to mind with attributes $(\$250, 4.2)$, then $s(A) = 0$, $s(B) = 1$, and B is chosen.

We now offer a fuller example to highlight the role of the inferred distribution. Consider options A and B in Figure 2. Each of the n comparison items increases the probability of choosing A, the probability of choosing B, or neither. The effect of a comparator will depend on where in the price $\times$ quality space it falls. Consider the regions marked R_A, R_B, and R_{AB}. R_B is the dominance region exclusive to option B, which we refer to as B's 'solo-dominance region'. Any item that falls within R_B is more expensive than both A and B, lower quality than B, and higher quality than A. Thus B dominates (i.e., is better than on both attributes) any item in R_B. Items in R_B are more expensive but also higher quality than option A, so they are not dominated by A. Similarly, comparison items that fall in R_A are dominated exclusively by A because they are only more expensive and lower quality than A. Finally, items that fall in R_{AB} are dominated by both A and B (i.e., in R_{AB} items are both more expensive and of lower quality than both A and B). MADS assumes that choice is determined by how many comparison items fall into each of these regions. The distribution from which comparison items are drawn will therefore affect whether A or B is chosen. The shaded area in Figure 2 represents such a distribution. A larger portion of the shaded area falls into R_A than R_B, meaning that A is more likely to be chosen.

Because of the probabilistic nature of the sampling process, A will not always be chosen, especially if the comparison sample is small, leading to a stochastic element of choice in our model.

Figure 2: Multi-Attribute Decision by Sampling

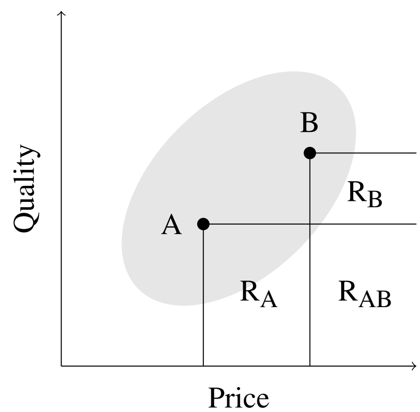

For ease of exposition, assume that the gray area corresponds to a uniform-height density which integrates to one. As lower prices and higher quality are preferred, a point is dominated if it lies to the south-east of another.

More generally, MADS assumes that an alternative accrues a point when an item falls in its solo-dominance region. For items in a joint-dominance region, a point accrues to any of the dominating alternatives' scores with equal probability. Therefore, if there are no dominance relations between the alternatives in the choice set, the alternative with the highest probability of accumulating a point is also the most likely to be chosen. The model provides analytic expressions for the probabilities of choosing each alternative from a choice set.[2] As such, given the choice set, the number of comparators n and the inferred distribution, one can compute the probability of each choice alternative being chosen.

[2]See Appendix B for details.

The Model

MADS describes two stages of processing. In the first stage, an sample is generated. Each alternative is $J-$dimensional, where $x = (x_1, ..., x_J)$ describes the level of each attribute of alternative x. Each $x \in X$ is referred to as a 'choice alternative', and X the 'choice set'. Given X, the individual then infers a $J-$dimensional distribution or 'inferred distribution' over the product space, with CDF denoted F_X. $n > 0$ draws are made from F_X. We assume draws are independently and identically distributed for the purposes of this paper. The set of draws sampled is denoted W where a typical element is $w = (w_1, ..., w_J) \in W$.

In the second stage, a score for each choice alternative is determined and a choice is made. The score of an alternative x is constructed via ordinal binary comparisons of the attribute levels against the other items in the 'reference set' $X \cup W \setminus x$. Elements of this set are referred to as 'comparators' or 'comparison items'. A choice alternative accrues a point when it is compared to an item in the reference set that it dominates. The choice alternative with the highest total score is then chosen. To represent this explicitly, let $\succsim_j$ be the rational binary preference relation which an individual has over levels of the attributes $j = 1, \ldots, J$ over any two items. Therefore x dominates y if $x \succsim_j y$ for $j = 1, \ldots, J$. In the case of hotels, where the attributes are price (p) and rating (q), both preference relations are assumed monotonic: $x \succsim_p y \iff x_p \leq y_p$ and $x \succsim_q y \iff x_q \geq y_q$. Given the reference set, the choice correspondence $c : X \mapsto X$ can be expressed as:

$$c(X) \in \underset{x \in X}{\arg\max} \{s(x)\}$$

where,

$$s(x) = |\{w \in W \cup X : A(w) = x\}|$$

$$A(w) = \begin{cases} \text{one member of } D_X(w) \text{, selected with probability } \frac{1}{|D_X(w)|} \text{ if } D_X(w) \neq \emptyset \\ \emptyset \text{ else} \end{cases}$$

$$D_X(w) = \{x \in X \setminus w : x \succsim_j w, j = 1, \ldots, J\}$$

If $c(X)$ is a singleton, then this item is chosen. If it contains more than one element, then each element of the set is chosen with equal probability. In this notation, $D_X(w)$ is the set of choice alternatives in X that dominate w (excluding comparison with itself), $A(w)$ is the alternative that accumulates a point from comparison item w, and $s(x)$ is the total score accumulated for each choice alternative.

Explaining Context Effects

We now provide an intuitive description of how the model captures each of the three context effects. As an accompaniment, we provide a statement of sufficient conditions on the inferred distributions for MADS to yield each context effect in Appendix B.

Attraction Effect

An illustration of how the inferred distribution may reasonably depend on the attraction effect choice sets is given by Figure 3. The triangular nature of the attraction effect choice sets pulls more of the density (shaded area) of the inferred distribution to the target's dominance region. This increases the probability that comparison items are drawn from the target's dominance region. This will tend to increase the score accumulated by the target relative to the non-target alternative, and hence the probability of it being chosen. Furthermore, attraction effect choice sets include a dominated alternative which gives the target alternative a head-start in the accumulation of points.

Figure 3: Attraction effect

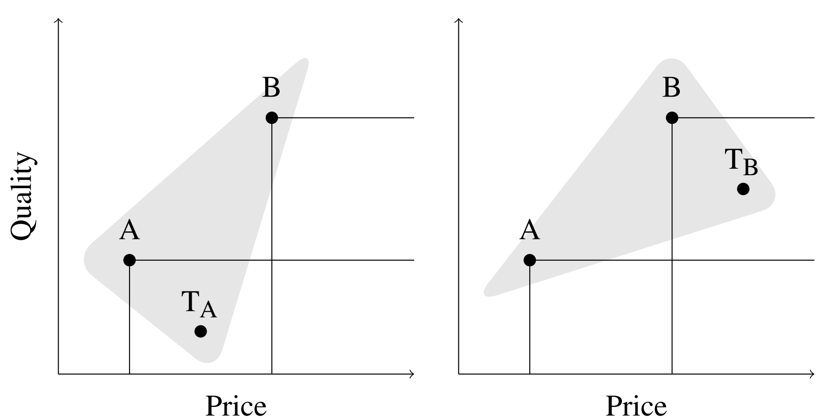

Compromise Effect

The intuition behind the explanation is illustrated by Figure 4. It is reasonable to suggest that the central parts of individuals' inferred distributions are wider. This ensures that more density is in the dominance region of the central (target) alternative. When the shift in the choice set is accompanied by a corresponding shift of the inferred distribution as shown, the central alternative remains the most popular: the compromise effect occurs. In fact, the compromise effect is the only one of the three for which our model requires a shift in the inferred distribution across choice sets.

Figure 4: Compromise effect

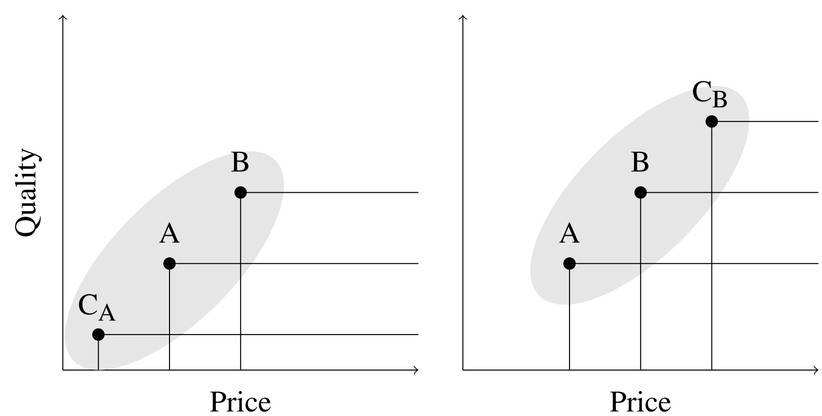

Similarity Effect

The intuition is illustrated by Figure 5. The similarity effect is driven by the fact that the non-target alternative is forced to share its solo-dominance region with the decoy. On the other hand, the target alternative is left with a relatively large solo-dominance region. When the decoy switches to be close to the other alternative, so that the non-target becomes the target, these forces reverse, causing the similarity effect. As the decoy becomes more similar to the non-target, this force is stronger, and the similarity effect larger. Conversely, as the decoy moves further away and towards the center, the choice sets will become the same (resembling a compromise choice set), and accordingly we would predict no difference in choice proportions.

Figure 5: Similarity effect

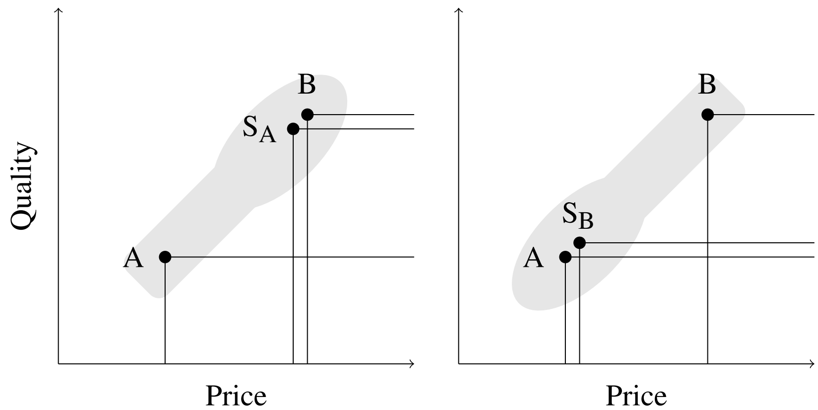

5. Experiments

We conducted two experiments to test MADS. Experiment 1 tested whether choice, and hence context effects, can be manipulated in accordance with our model's predictions. Experiment 2 tests the prediction that the inferred distribution will depend on the choice set and allows us to estimate the parameter n.

Hotel Price and Quality Data

Choice alternatives for the experiments were Manhattan hotel stays, for which there are two main attributes: 'average rating' and 'price'. Data pertaining to real hotels were taken from Hotels.com on 23 June 2014. We recorded the price and average rating of the cheapest 200 hotel stays for a one-night stay for one adult in one room, for a stay on 12 November 2014.[3] 'Average rating' refers to the rating given by members of Hotels.com who had previously stayed at the hotel.[4] We presented the score rounded to one decimal place, as it is presented on the website itself. This served as our proxy for quality, so that we could present data across the price-quality domain, as in classical context-effect experiments. Given the familiarity of such sites to internet users, we referred to 'average rating' rather than 'quality' throughout the experiment. 12 distinct hotels were selected from these data to form the six choice sets of the three context effects.[5] As with most

[3] See Figure A1 in Appendix A for a plot of the hotel data recorded.

[4] Each reviewer submits a score of 1, 2, 3, 4 or 5, where higher numbers correspond to a better experience. All the hotels we recorded had at least 25 reviews.

[5] Details of the choice sets used can be found in Table A1.

studies showing the presence of these context effects, participants' hotel choices were hypothetical.

Study 1: Choice Manipulation

Method

1,304 Amazon Mechanical Turk (AMT) workers were recruited in July 2014.[6] 68 were excluded from the analysis because they had previously completed a related pilot study; one was removed because they did not complete the experiment. This left data from 1,235 participants for analysis. Average completion time was 14 minutes 27 seconds. Participants were compensated with a participation fee of $1.50, which corresponds to an average hourly wage of $6.23. Participants were randomly assigned to one of 12 conditions $= 3$ [attraction, compromise, similarity] $\times 2$ [A target, B target] $= 6$ choice sets $\times 2$ conditions [treatment, control].

Participants in the treatment condition each saw data relating to ten hotels taken from the dataset before selecting an alternative from one of the six choice sets. They were shown the price and average rating of each of these ten hotels one at a time, and for each one were asked to indicate on a seven-point scale how likely they would be to stay at the hotel. The purpose of asking this was to ensure some amount of engagement with the hotels presented, such that they would affect the hotels available in the participants' comparison sample. An example screen-shot is provided in Figure 6. Following the treatment, participants faced one of the choice sets and answered the question "Which hotel would you be most likely to choose?". Participants in the control condition simply chose without seeing any other hotels first. Before finishing, participants faced a series of questions for another experiment. Participants were then asked to indicate how they divided their attention when considering hotel choices using a seven-point scale where 1 meant "considered solely prices", 4 meant "both attributes equally", and 7 "solely ratings". Basic demographic questions followed on the final screen.

[6]The rationale for this sample size is given in Appendix A

Figure 6: Screen-shot of manipulation

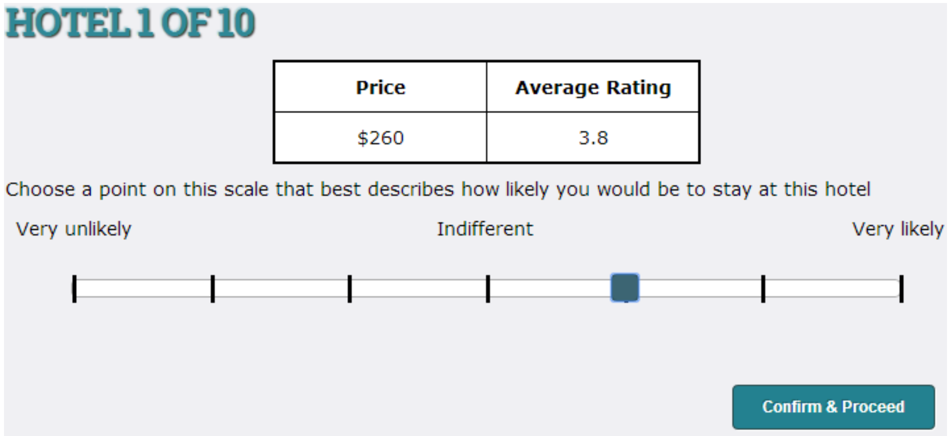

Note that the slider was not visible until the participant clicked on the scale.

The ten hotels shown as a treatment were chosen such that they were dominated by alternative B (the more expensive, higher-quality alternative), but not dominated by alternative A (the cheaper, lower-quality alternative), in the choice set they faced afterwards. Because the hotels in each context effect choice set were different, a different set of hotels were used as a manipulation for each choice set. Where there were more than ten candidate hotels fitting this description, ten were chosen at random. Every participant in the same choice set saw the same ten hotels, but in a random order. The manipulation is described by Figure 7. Notice that in the attraction and compromise choice sets, only alternative B is promoted. In the similarity set including a decoy close to B, both B and the decoy are promoted, because they are close together. We now refer to the target alternatives in general as the 'manipulation targets'.

Figure 7: Manipulating choice: method

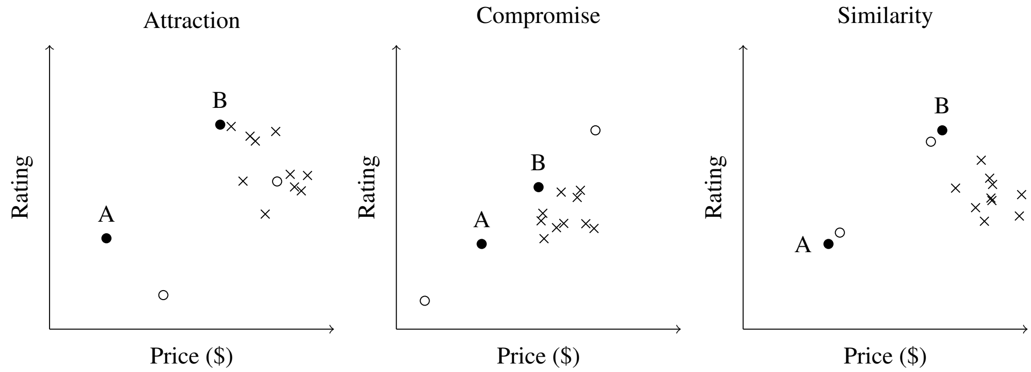

Participants saw ten hotels (crosses) in a random order prior to facing a ternary choice set: A and B plus one of the decoys (unfilled circles). Notice that in the attraction and compromise choice sets, only alternative B is promoted. In the similarity set including a decoy close to B, both B and the decoy are promoted, because they are so close together.

Results

The treatment effect is the difference in the proportion of participants choosing the manipulation targets in the treatment and control conditions. Overall, participants chose the manipulation targets 7% more often as a result of the treatment ($p < .001$).

Our theory supposes that this manipulation will be successful because the hotels that participants were shown are dominated by the manipulation targets on both attributes, but by the other alternatives in the choice set on only one attribute. Therefore, we expect to the manipulation to have the most impact when individuals pay attention to both attributes. Table 1 shows large differences in manipulation effect depending on whether or not participants paid attention to both attributes equally. The majority of participants paid attention to both attributes equally. The effect of the treatment on these participants was 17% ($p < .001$), but was zero for those reporting that they paid attention unequally.

Table 1: Manipulating choice.

Attribute attention	Observations	χ^2 statistic	Manipulation effect
All (1-7)	1,235	9,397,860	.07***
Equal (4)	632	1,286,446	.17***
Non-equal (1-3, 5-7)	603	0.17	.01

Attribute attention: sliding scale of integers $\{1, ..., 7\}$ where 1=only considered price, 4=considered both attributes equally, 7=only considered quality. Manipulation effect: the difference in the percentage of participants choosing the manipulation targets in the treatment and control conditions. *** denotes a significant difference from zero, where $p < .001$ from χ^2 tests against the null of no effect. No star denotes $p = .681$.

Our theory is focused on dominance on all attributes. The manipulations had a significant and large impact on the choice for those paying equal attention to both attributes. For those not doing so, no effect was found. We therefore continue the analysis using the data from those who reported paying equal attention.[7]

Our design permits us to attempt to counter and enhance the three context effects. For example, the attraction effect says that alternative B will be chosen more often from $\{A, B, T_B\}$ than $\{A, B, T_A\}$, with no manipulation. When participants choosing from $\{A, B, T_A\}$ are instead in the treatment condition, we predict that B will be more popular than if they were not. Therefore, when we compare choices from participants who faced $\{A, B, T_B\}$ in the control and $\{A, B, T_A\}$ in the treatment group, we predict that the attraction effect will be reduced. This example corresponds to the first column in Figure 8. The remaining columns describe which data from which conditions are compared to investigate the impact on the context effects.

[7]See Figure A2 for the results including all participants.

Figure 8: Manipulating context effects: method.

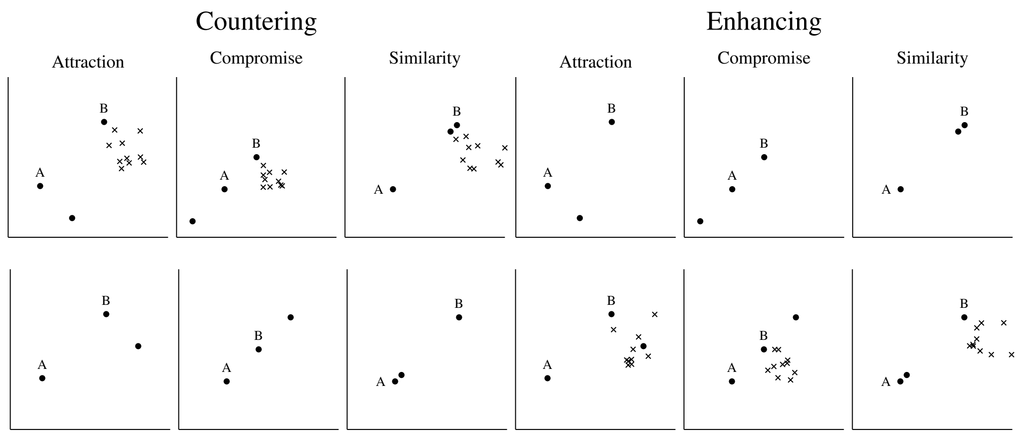

In the left (right) three columns, we predict that context effects will be countered (enhanced) due to the manipulation. Choice sets are shown as solid circles and the hotels of the manipulation as crosses. Price is on the x-axis, average rating on the y-axis.

We know from Table 1 that the manipulations had a significant impact on choices. The manipulation effect is still significant when broken down by context effect (all three χ^2 tests give $p < .001$). Figure 9 shows how choice manipulations in turn affected the presence and strength of the attraction, compromise and similarity effects.

Figure 9: Manipulating context effects: evidence.

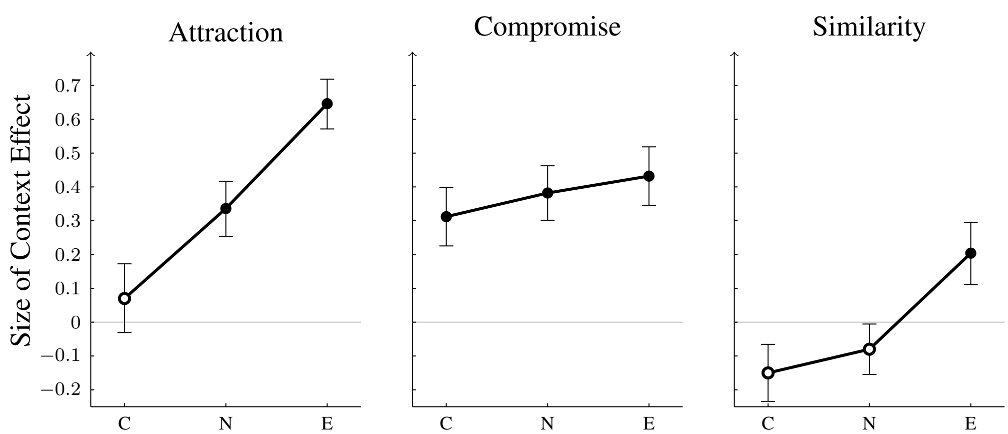

Context effect size is the difference in the choice proportions of alternative B. (C) denotes countering and (E) enhancing where the choice data used for each are depicted in the columns of Figure 8; (N) denotes ‘neutral’, which refers to the standard attempt to replicate the context effects. Standard error bars are given. Solid circles refer to a significant context effect i.e., different from zero, at the 5% level from a standard χ^2 test. Hollow circles refer to no significant difference from zero, and hence no context effect.

The attraction effect was replicated with an effect of .34, significantly different from zero.[8, 9] When we enhanced the context effect through our manipulation, the effect size almost doubled, to .65. When we countered the effect, the effect fell to .07, which is insignificantly different from zero. In other words, the attraction effect was turned off. The compromise effect was .38, significantly different from zero. Countering the compromise effect reduced it to .31 and enhancing lifted it to .43.

The similarity effect was not replicated, with an effect insignificantly different from zero. Our theory dictates that the presence of context effects are probabilistic and we are not the first study to find that the similarity effect is the weakest (e.g., Noguchi and Stewart, 2014).[10] More importantly, our manipulations did have an effect. Countering the similarity effect pushed the size down to -.15, which is marginally significant ($p =$ .07), which shows some evidence of reversing the effect. When we enhanced the effect, the size became .20, which is significantly different from zero. In other words, the similarity effect was turned on.

Study 2: Distribution Elicitation

Method

607 Amazon Mechanical Turk workers were recruited. The participants of study 2 were also the control group for study 1. Following their choice, and hence exposure to a choice set, we elicited what participants inferred about the rest of the hotel market. It would have been prohibitive to ask them for their best guess of all the other 197 hotels in our dataset. Instead we asked them to estimate the price and quality of a randomly chosen 20. They were told that 20 hotels had been randomly selected from our dataset, which they had to estimate. These 20 did not include the three they had already seen in their choice set. Participants were able to see their choice set throughout the elicitation

[8] An effect of .34 means that 34% more people chose B from $\{A, B, T_B\}$ than from $\{A, B, T_A\}$.

[9] Recent work has investigated and debated the domains in which the attraction effect can be expected to appear (Frederick *et al.*, 2014; Huber *et al.*, 2014; Simonson, 2014; Yang and Lynn, 2014). This work suggests that pictorial and verbal (rather than numerical) descriptions of choice alternatives can dampen or nullify the attraction effect. Although it is important to recognize the boundaries under which an effect is expected, quality (in addition to price) is often represented by a numerical value, especially in the modern online marketplace.

[10] Note that when we looked at individuals' inferred distributions we did not significantly predict choice in the similarity effect conditions, but did in the attraction and compromise effect conditions (see the estimation section).

process, but not to change their choice.

To elicit their estimation of the market distribution, they completed two screens. On the second screen, we asked participants to plot where they thought these 20 hotels lay in price$\times$rating space, based on the choice set they had seen. That is, they were shown their choice set plotted on a pair of axes and required to plot where they thought the 20 additional hotels were located. An example screen-capture is provided in Figure 10. The example participant shown chose from a similarity effect choice set, $\{A, B, S_A\}$, which typically promotes the cheapest option A (labeled here for the benefit of the reader). The participant has placed nine plots so far. The choice set was displayed in red; plots in green. To avoid possible anchoring effects, we did not include grid lines, axis ticks or axis-tick labels. To scale the axes, we asked for their best guess of the minimum and maximum price and average rating of the 23 hotels (including the choice set they faced). The minimum and maximum value of each axis was determined by participants themselves on the first screen, which they could not return to. Illogical answers e.g., that the minimum was higher than the maximum, were not allowed. The on-screen size of the plotter was fixed to be a square; participants only determined scale of the axes. Participants were able to see the coordinates where their mouse was hovering, were able to reset the graph and were provided with a counter telling them how many points they still had to place.

Figure 10: Screen-capture of distribution elicitation.

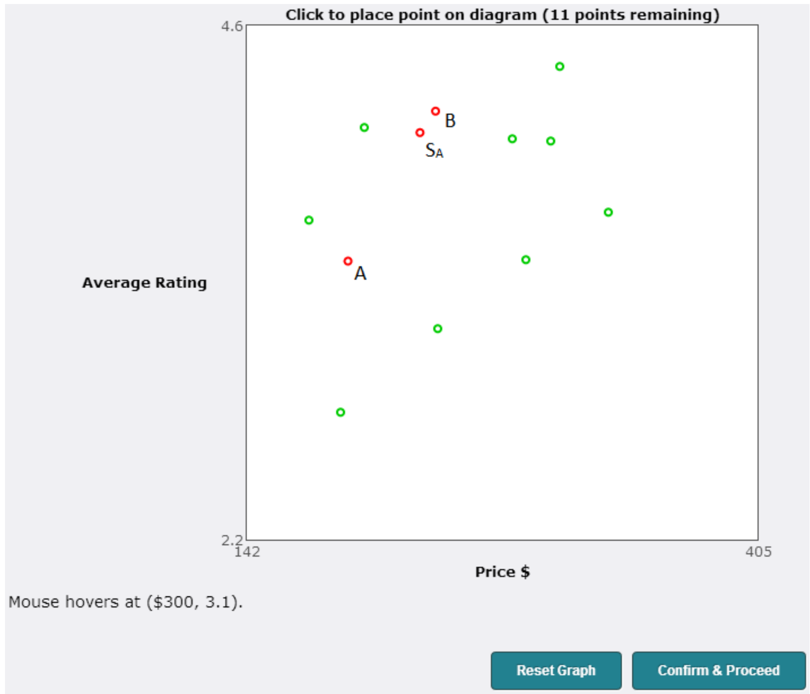

We did not allow participants to go back to change the minimum and maximum values for the axes due to concerns that participants would tweak their answers to move their choice set around the plotter. We removed data of participants who entered values extreme enough such that the choice set would be shown bunched into the corner of the screen.[11] We chose the cutoff to be a price of $800; anyone entering this value or higher was excluded. This removed 35 participants, leaving 572 for analysis.

As an incentive payment, participants were told that the five who plotted closest to the 20 hotels would be paid $5 as a bonus. The procedure we used to determine who was the closest was that according to the modified-Hausdorff metric, as advocated by Dubuisson and Jain (1994). This metric provides a distance based on the average minimum pairwise distances between two sets of coordinates. In our case, these two sets were the participant's plotted data, and the 20 hotels randomly chosen.[12]

[11]Confirming this concern, we were in fact contacted by one participant who had entered a maximum price of $5,000, who wanted to explain his choice to us.

[12]Participants were not told the details of the metric; they were simply told that the five plots 'closest' to the 20 we had would be paid.

Results

Distribution Elicitation

Each panel of Figure 11 presents the pooled plots placed by all ≈100 participants per choice set, meaning there are roughly 2,000 plots in each panel. We provide the proportions of plots contained within the crucial areas identified by theory. We draw on patterns in these aggregate data which illustrate how choice sets affect distributions, and that these distributions move in ways compatible with our theory to produce the context effects. We discuss features of this aggregate distribution data as if it were in fact the distribution of all individuals in order to demonstrate how we consider context effects can arise. Recall that MADS supposes that movements in the inferred distributions change the probabilities of alternatives being included in the comparison set, and hence affect choice probabilities. Specifically, if there are no dominance relations between the alternatives in the choice set, the alternative with the highest probability of accumulating a point is also the most likely to be chosen.

Figure 11: Plotting data by choice set and aggregate-PDF values.

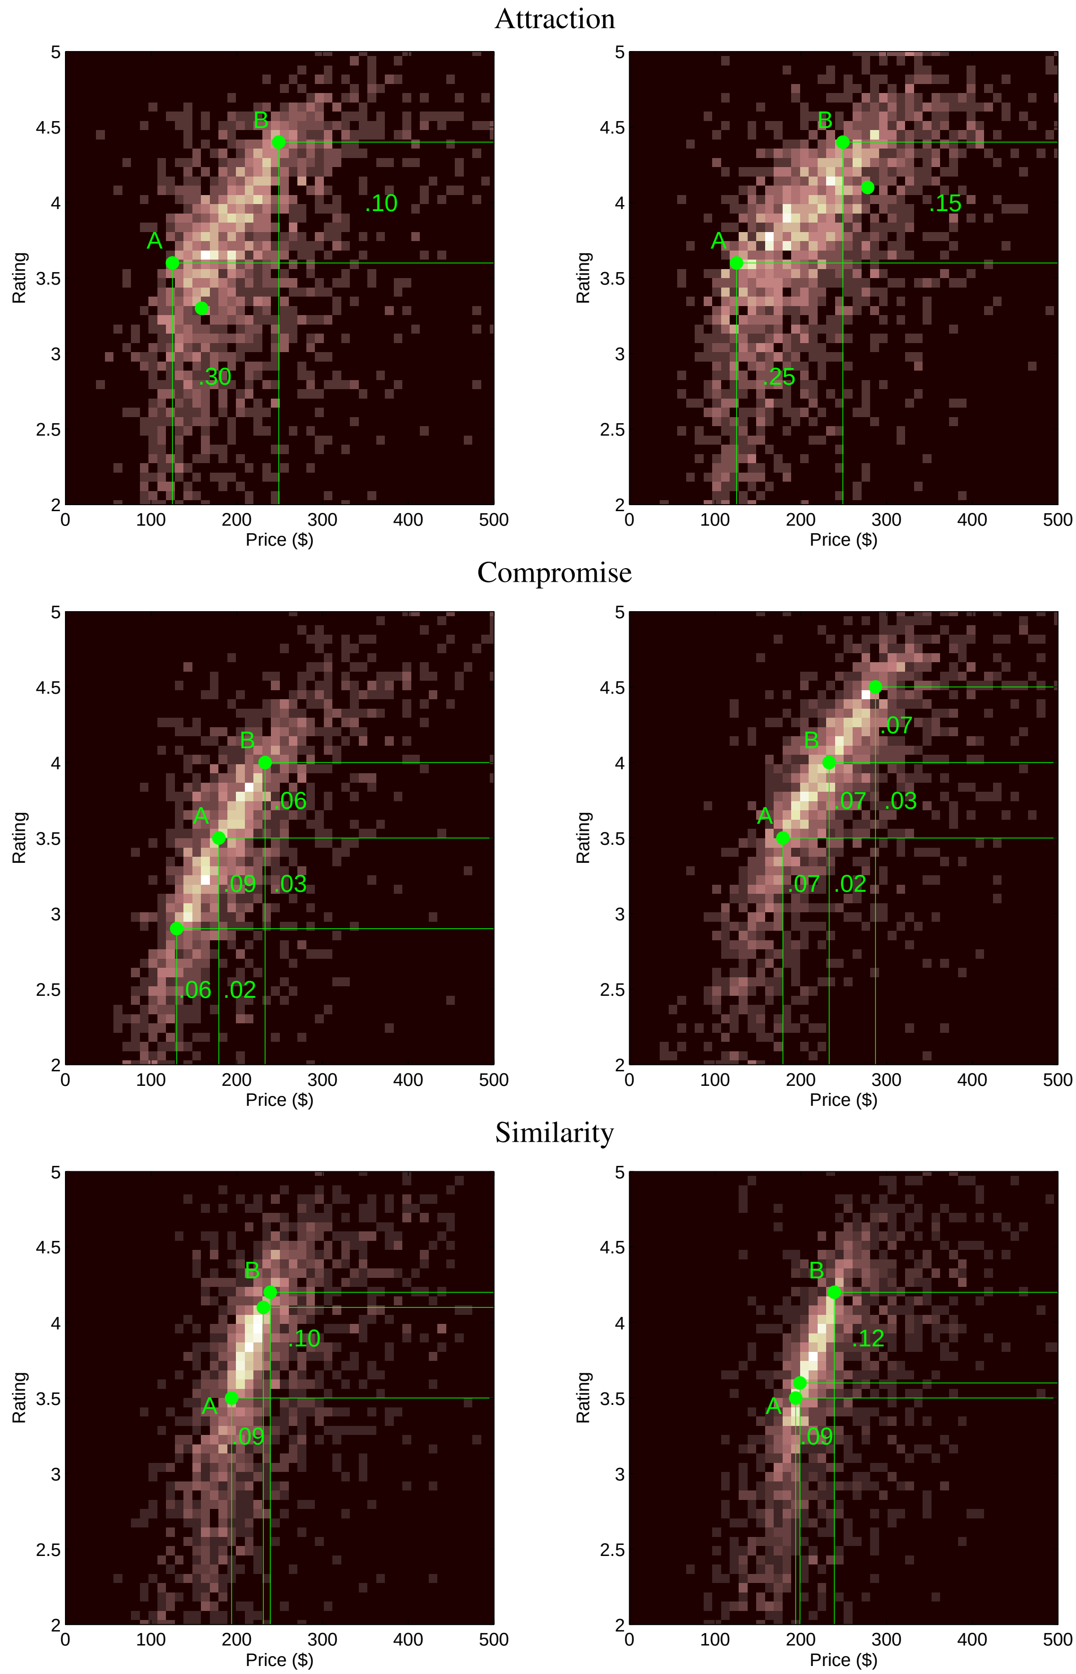

Lighter colors refer to a higher density of plots. The numbers refer to the proportion of points plotted in that region i.e., the empirical density. Graphics are cropped at 2/5 for quality, and $500; over 95% of the plotting data is in this range.

Inspection of the heat-maps reveals pronounced movement of the density of the plotted points between choice sets. Using a multi-dimensional version of the Kolmogorov-Smirnov test, as proposed by Fasano and Franceschini (1987) we found that within each of the three context effects the distributions elicited from participants differed significantly ($p < .001$) depending on which choice set they had seen.

The attraction effect panels show that the total density of the solo-dominance regions is .40 in both conditions. .05 of the density shifts from A's dominance region when A is the target to B's dominance region when B is the target. The shift in density towards the target, and the fact that the target enjoys a 1-0 head-start in the score accumulation from the decoy, combine to predict the attraction effect.

In the compromise effect panels, .06 of the density is in B's dominance region under $\{A, B, C_A\}$ and $\frac{.03}{2} = .015$ in the region dominated by both B and the decoy. This gives the probability of B accumulating a point as approximately $.075$. When B becomes the target, the probability of B accumulating a point is approximately $.095$, a 27% increase. For A, the corresponding numbers are $.08$ and $.095$, a 19% increase when it is the target. The movements in density are consistent with the MADS account of the compromise effect.

In the similarity effect panels, B has a negligibly small solo-dominance region because it is similar to the decoy and .10 of the density is dominated by both. The narrow regions all have less than .03 density and are ignored for this discussion. Accordingly, B has approximately .05 probability of accumulating a point. When B becomes the target, the same area of price$\times$quality space is dominated only by B. The probability of B accumulating a point under $\{A, B, S_B\}$ is approximately .12. This illustrates how MADS explains the similarity effect.

Estimation

Finally, despite the small sample size per participant, we examined whether it was possible to predict individuals' choices. Our model imposes no assumptions or parameters on an individual's inferred distribution. Only 20 plots per participant are available, so individual empirical distributions are too coarse to enable prediction. Therefore, we selected the multi-variate distribution that best-fitted each participant's plots (i.e., selection was independent of the choice data). Various copulas (Gaussian, t, Frank, Gumbel and Clayton) were fitted to each participant's estimate of the distribution of 23 points (20 plotted and 3 from their choice set). (Copulas are succinct descriptions of the correlation between two variables.) Copula selection for each participant was determined by the Akaike Information Criterion.

For each participant we determined the choice that was most likely on the basis of that participant's estimated inferred distribution. The proportion of correct predictions is reported in Table 2. It can be seen that the choices were reasonably well predicted for participants in the attraction and compromise conditions but not in the similarity condition. These findings are congruent with the fact that we replicated the attraction and compromise effects but not the similarity effect, as shown in Figure 9.

Table 2: Predicting individuals' choices.

Context effect	Observations	Null	Predicted	p-value
Attraction	194	.50	.62	.001
Compromise	180	$.\overline{33}$	.42	.014
Similarity	198	$.\overline{33}$	.30	.327

The null hypotheses are those implied by random prediction between all non-dominated choice alternatives. p-values are from two-sided binomial tests.

To compute the estimates of Table 2, we first estimated the MADS parameter n that specifies the number of comparison items brought to mind from individuals' inferred distributions. We calculated the probability of participants' choice data for different values of n and found the maximum likelihood estimate to be 4.[13] The estimate is precise in the sense that it is significantly different from both 3 and 5 (LR tests $p = .021$ and $p = .091$ respectively). Note that $n = 4$ is a psychologically realistic value for

[13]Three participants chose the decoy alternative from attraction effect choice sets and were excluded from the estimation, leaving 569 participants. Our model has no additional error term and so predicts such behavior with probability zero, which prevents the likelihood from being well-defined.

working memory capacity (see Cowan, 2001), consistent with the idea that comparison samples are held in working memory during choice.

6. Discussion

MADS offers a single-mechanism model of the attraction, compromise and similarity effects. In Experiment 1, prior exposure to a distribution of market options altered subsequent choices in ways predicted by the model, and allowed us to remove and enhance the classic context effects. This result, consistent with recent studies of risky choice (e.g., Stewart *et al.*, 2015, Ungemach *et al.*, 2011) but here pertaining to the "big three" context effects, suggests that assumptions about background distributions, combined with choices made on the basis of simple dominance relations, may give rise to context effects. Experiment 2 demonstrated that individuals' inferred distributions of market options depend on the choice set in ways required by the model, thus extending and applying previous studies within the "options as information" tradition (e.g., Sher and McKenzie, 2014, Shenoy and Yu, 2013, Kamenica, 2008). Moreover, the maximum likelihood estimator of the model's central parameter took on a psychologically realistic value for the capacity of human working memory.

Our data was generated by members of Amazon's online platform, Mechanical Turk (AMT). AMT's online population has been shown to have the advantages of being more demographically diverse, and producing data of the same, if not better, quality than more traditional participant methods (Paolacci and Chandler, 2014; Chandler *et al.*, 2014). This has been shown through many studies replicating classic experiments in various domains including cognitive psychology (e.g., Paolacci *et al.*, 2010, Goodman *et al.*, 2013) and economics (e.g., Horton *et al.*, 2011). Some have expressed concern that AMT participants may not pay sufficient attention to the choice alternatives which may lead to a failure to find the attraction effect (Simonson, 2014). We demonstrate that the classic context effects in consumer choice can be found with such samples, and are the first to our knowledge to demonstrate the presence of all three (where the similarity effect was found only after the treatment).

Our modeling approach contrasts with existing models of these context effects in that it does not assume value-like encoding of stimulus attributes (albeit values modified by attentional weights, availability, etc). Vlaev *et al.* (2011) distinguish between three broad classes of model: those that assume the computation of different options' values (which then inform choice directly); those that assume value computation modified by context, and "comparison-only" models in which choices depend directly on (typically ordinal-only) comparisons without involving any computation of value. Our model is an example of the third, comparison-only, category of model.

More specifically, it can be seen as an extension of rank-based models such as Decision by Sampling (Stewart *et al.*, 2006) and, more generally, sampling models that assume that options are evaluated relative to an assumed background distribution (see also Kornienko, 2013). According to DbS, judgments result from a series of binary ordinal comparisons of a to-be-judged target against other values in a mental sample. DbS emphasizes the role of comparison items retrieved from memory. For example, consider evaluating a hotel price of \$200 per night. DbS suggests that a sample of other hotel prices would be brought to mind. These could be remembered prices paid in the past, and/or other prices visible on a hotel comparison website. If the mental sample contained two hotel prices below \$200, and five higher prices, the subjective judgment of the \$200 price would be determined by its relative rank position in the context, i.e., $\frac{5}{7}$. An implication is that different mental samples (reflecting differences in individuals beliefs about hotel prices elsewhere) would lead to different judgments about the same actual price. MADS both extends DbS to the two-dimensional case (see also Stewart and Simpson, 2008), and specifies the role of choice options in causing the background distribution to be updated.

In economics, the classical economic decision-maker is assumed to be able to effortlessly make trade-offs across multiple dimensions. Recently, work has focused on equilibrium in markets when consumers exhibit some degree of trade-off aversion, employing some heuristic instead (e.g., Bachi and Spiegler, 2015; Papi, 2014). In particular, Bachi and Spiegler (2015) study two-attribute goods and assume that a dominant alternative is chosen when it exists.

More generally, much experimental evidence is consistent with the idea that the ranked position of an item within an experimental context will influence its judgment. Evidence came initially from psychophysics (Parducci *et al.*, 1960; Parducci and Perrett, 1971) but has since been found in subjective judgments of sweetness (Riskey *et al.*, 1979), morality (Marsh and Parducci, 1978), body image (Wedell *et al.*, 2005), attractiveness (Wedell *et al.*, 1987), fairness (Mellers, 1982), personality (Wood *et al.*, 2012), and prices (Niedrich *et al.*, 2001; Niedrich *et al.*, 2009). More recently, rank-based and range-frequency theory principles have been used to understand sensitivity to human fatalities (Olivola and Sagara, 2009), individuals' satisfaction with both their wages (Brown *et al.*, 2008) and their life in general (Boyce *et al.*, 2010; Smith *et al.*, 1989).

Our approach is therefore consistent with range of evidence for rank-based processes in various areas of psychology. Moreover, individual differences in beliefs about background distributions influence judgment. For example, individuals' beliefs about whether they are drinking too much or exercising enough depend not just on their objective behavior (e.g., the number of days per week they exercise; the amount of alcohol they consume) but on how they believe their symptoms rank within the general population (i.e., how many people out of 100 they believe drink or exercise more than they do) (Maltby *et al.*, 2012; Wood *et al.*, 2012). Beliefs about the relevant background distributions also predict people's judgments of whether they are depressed or anxious (Melrose *et al.*, 2013), students' satisfaction with their educational experience (Brown *et al.*, 2015b) and attitudes towards indebtedness (Aldrovandi *et al.*, 2014), and individuals' attitudes towards dishonest behavior (Aldrovandi *et al.*, 2013). In summary, MADS not only provides a one-parameter account of the "big three" context effects; it is also consistent with a wide range of psychological data from other domains.

Appendix A - Methodological Details

Sample Size

It was decided in advance that 1,200 participants would be tested; 1,300 were requested from AMT in order to be able to remove some if there were those who had completed a related pilot, and 1,304 were received. There is considerable variation in the size of context effects in the literature, and they are of course not always found. Estimates from studies in consumer choice find sizes ranging from about 0.15 (see Table 1 of Trueblood *et al.* 2013) to over 0.3 (Noguchi and Stewart 2014). We chose to ensure 100 participants per condition; this gives a power of 0.8 to detect a difference in choice proportion of 0.2 when comparing two conditions with each other.

Supporting Figures

Figure A1: Hotel data.

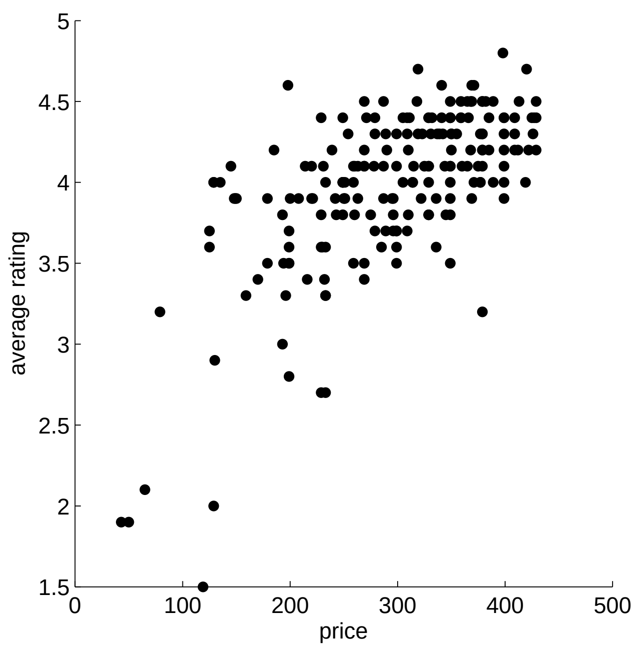

Table A1: Choice sets (price, rating).

Alternative	Attraction	Compromise	Similarity
A	125 3.6	179 3.5	194 3.5
B	249 4.4	233 4.0	239 4.2
D_A	159 3.3	130 2.9	231 4.1
D_B	278 4.1	287 4.5	199 3.6

Figure A2: Manipulating context effects: evidence (all participants).

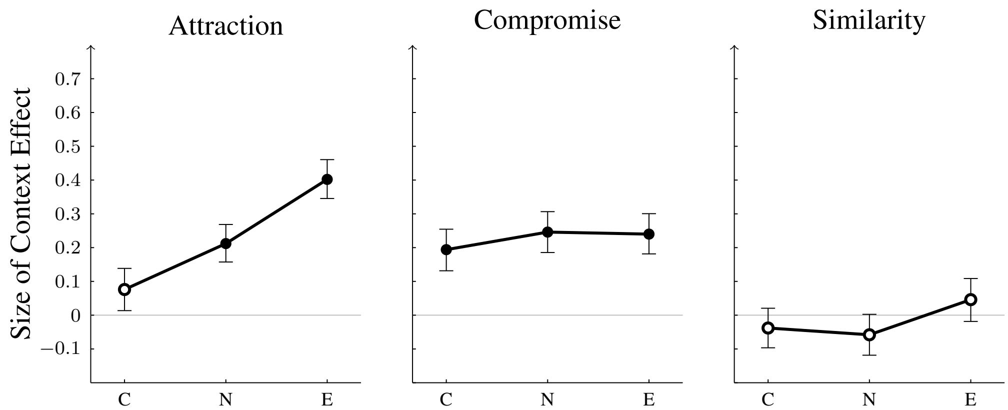

The same notes as Figure 9 apply.

Appendix B - Sufficient Conditions for Context Effects

We now provide sufficient conditions on the inferred distributions for MADS to predict each of the three context effects. Because the expressions for choice probabilities are lengthy, we provide sufficient conditions under 'symmetry' which provide a cleaner statement for how the model can generate the effects. Figs. 3-5 from the main text are counterparts to the Propositions which display distributions that satisfy symmetry. Proofs are detailed below.

Denote the probability of an inferred comparison item being dominated by an alternative x having been presented with choice set X, as $F_X(x)$. Furthermore, let $F_{\{A,B,D_A\}}$ and $F_{\{A,B,D_B\}}$ be denoted by $\mathcal{A}$ and $\mathcal{B}$ respectively, where D_A and D_B are the decoys making A and B the targets respectively, where the exact positioning of the decoy depends on which context effect is in question. This means for example, that $\mathcal{A}(B)$ refers to the probability of drawing an item which is dominated by B (and not A) from the inferred distribution induced by $\{A, B, D_A\}$. For joint dominance regions, we analogously denote $\mathcal{A}(A, B)$ ($\mathcal{B}(A, B)$) as the probability of drawing a n item which is dominated by both A and B from the inferred distribution induced by $\{A, B, D_A\}$ ($\{A, B, D_B\}$).

Attraction Effect

Proposition 1 (Attraction)**.** *The model produces the attraction effect if:*

(i) Symmetry: $\mathcal{A}(A, B) = \mathcal{B}(A, B)$, $\mathcal{A}(A) = \mathcal{B}(B)$ *and* $\mathcal{A}(B) = \mathcal{B}(A)$.

(ii) $\dfrac{\mathcal{A}(A)}{\mathcal{A}(B)} \left[= \dfrac{\mathcal{B}(B)}{\mathcal{B}(A)}\right] > 1.$

The first symmetry condition assumes that the probability of sampling an item which affects the expected scores of A and B by the same amount, is the same. If this is satisfied, one can focus on the probabilities of items falling in regions that affect the *difference* in the scores. The second and third equalities in (i) impose symmetry on the inferred distributions around A and B, so the explanation can abstract from differences that may occur because A is cheaper, or B is higher quality. The final statement drives our intuition. It states that the probability of an item falling in the solo-dominance region of the target is greater than of an item falling in the solo-dominance region of the non-target alternative. This implies it is more likely that the target accumulates a point. Because the item with the highest score is chosen, these conditions produce the attraction effect.

Compromise Effect

Proposition 2 (Compromise)**.** *The model produces the compromise effect if the following are satisfied:*

(i) Symmetry: $\mathcal{A}(A) = \mathcal{B}(B)$, $\mathcal{A}(B) = \mathcal{A}(C_A) = \mathcal{B}(A) = \mathcal{B}(C_B)$, $\mathcal{A}(A,B) = \mathcal{A}(A,C_A) = \mathcal{B}(A,B) = \mathcal{B}(B,C_B)$ *and* $\mathcal{A}(A,B,C_A) = \mathcal{B}(A,B,C_B)$.

(ii) $$\frac{\mathcal{A}(A) + \frac{1}{2}\mathcal{A}(A,B) + \frac{1}{2}\mathcal{A}(A,C_A)}{\mathcal{B}(A) + \frac{1}{2}\mathcal{B}(A,B)} \left[= \frac{\mathcal{B}(B) + \frac{1}{2}\mathcal{B}(A,B) + \frac{1}{2}\mathcal{B}(B,C_B)}{\mathcal{A}(B) + \frac{1}{2}\mathcal{A}(A,B)} \right] > 1.$$

Condition (ii) drives the intuition about what is required for the compromise effect. Due to symmetry, the discussion is the same for both alternatives, so we will just refer to B. Under symmetry, the condition $\mathcal{B}(B) + \frac{1}{2}\mathcal{B}(A,B) + \frac{1}{2}\mathcal{B}(B,C_B) > \mathcal{A}(B) + \frac{1}{2}\mathcal{A}(A,B)$ implies that the probability of a point being accumulated to B's score is higher under $\{A,B,C_B\}$ than $\{A,B,C_A\}$. Hence under symmetry, (ii) is sufficient to yield $p(B|ABC_B) > p(B|ABC_A)$, the compromise effect. Note that if B's solo-dominance region is larger when it is the target it helps to produce the effect. Notice also, that when B is the target, it has two joint-dominance regions, but only one when it is not the target. This accounts for the presence of an extra joint-dominance region in the numerator of (ii).

Similarity Effect

Proposition 3 (Similarity)**.** *The model produces the similarity effect if the following are satisfied:*

(i) Symmetry: $\mathcal{A}(A) = \mathcal{B}(B)$, $\mathcal{A}(B) = \mathcal{B}(A)$, $\mathcal{A}(S_A) = \mathcal{B}(S_B)$, $\mathcal{A}(A,S_A) = \mathcal{B}(B,S_B)$, $\mathcal{A}(B,S_A) = \mathcal{B}(A,S_B)$ *and* $\mathcal{A}(A,B,S_A) = \mathcal{B}(A,B,S_B)$.

(ii) $\dfrac{\mathcal{A}(A)+\frac{1}{2}\mathcal{A}(A,S_A)}{\mathcal{B}(A)+\frac{1}{2}\mathcal{B}(A,S_B)}\left[=\dfrac{\mathcal{B}(B)+\frac{1}{2}\mathcal{B}(B,S_B)}{\mathcal{A}(B)+\frac{1}{2}\mathcal{A}(B,S_A)}\right]>1.$

Condition (ii) carries the intuition of what is required to produce the similarity effect. Under symmetry, the condition $\mathcal{B}(B)+\frac{1}{2}\mathcal{B}(B,S_B)>\mathcal{A}(B)+\frac{1}{2}\mathcal{A}(B,S_A)$ implies that the probability of B accruing a point is higher under $\{A,B,S_B\}$ than $\{A,B,S_A\}$. Hence under symmetry, (ii) is sufficient to yield $p(B|ABS_B)>p(B|ABS_A)$, the similarity effect. By inspection of Figure 5 one can see that $\mathcal{B}(B)$ is likely to be greater than $\mathcal{A}(B)$ due to the configuration of the similarity effect choice sets. Although $\mathcal{A}(B,S_A)$ is also likely to be greater than $\mathcal{B}(B,S_B)$, these joint-dominance regions only have a probability of $\frac{1}{2}$ of adding to B's score.

Symmetry

The symmetry conditions of Propositions 1-3 suppose that individuals' inferred distributions will depend on the relative position of the choice set's alternatives to each other, but will otherwise be the same i.e., will not depend on the absolute levels of the attributes. Symmetry allows for a clear exposition, allowing explanations to rely on a ratio consisting of a few areas of the inferred distribution's density. We consider it a natural benchmark case, especially for markets where individuals have had no prior experience. In practice, when individuals are assessing items, they will draw not only on the choice set presented to them, but also on their prior experience or knowledge of the product. This will also affect their inferred distribution over the product space. For hotels, if individuals have predominantly had exposure to cheaper, lower quality hotels than those in the choice sets offered, their inferred distributions will place more weight on this end of the market, in addition to any inference from the choice sets. This would cause the distributions to be skewed.

More general conditions which allow for a relaxation of symmetry are still expressible analytically, but no longer have the simplicity of those in Propositions 1-3, requiring conditions quantifying the asymmetry and on n. Where data does not satisfy symmetry, the full expressions for choice probabilities can be used directly to check when the context effects are expected.

Proofs

Recall $\mathcal{A}$ and $\mathcal{B}$ denote the inferred distributions given choice sets $F_{\{A,B,T_A\}}$ and $F_{\{A,B,T_B\}}$ respectively. Also recall that we have assumed i.i.d. draws from inferred distributions. The probability of choosing x from X is denoted $p(x|X)$. In equations, choice sets are written without braces or commas to save on the width of the text. To simplify the expressions, instead of using the inferred distribution directly, we use the probability of a choice alternative accumulating a point relative to other choice alternatives. These are denoted $\alpha(x)$ and $\beta(x)$ given choice sets $\{A, B, D_A\}$ and $\{A, B, D_B\}$ respectively. $n > 0$ is assumed, but in some instances $n = 1$ would cause a summand to be invalid, in those cases, the sum is zero.

Proof of Proposition 1:

$$\alpha(A) = \mathcal{A}(A) + \tfrac{1}{2}\mathcal{A}(A,B) \qquad \beta(A) = \mathcal{B}(A) + \tfrac{1}{2}\mathcal{B}(A,B)$$

$$\alpha(B) = \mathcal{A}(B) + \tfrac{1}{2}\mathcal{A}(A,B) \qquad \beta(B) = \mathcal{B}(B) + \tfrac{1}{2}\mathcal{B}(A,B)$$

$$\alpha_0 = 1 - \alpha(A) - \alpha(B) \qquad \beta_0 = 1 - \beta(A) - \beta(B)$$

These allow us to write:

$$\begin{aligned}
p(A|ABT_A) =& \frac{1}{2}\sum_{k=0}^{\lfloor\frac{n-1}{2}\rfloor}\binom{n}{k, k+1, n-2k-1}\alpha(A)^k\alpha(B)^{k+1}(1-\alpha(A)-\alpha(B))^{n-2k-1}\\
&+\sum_{k_1=0}^{n}\sum_{k_2=0}^{\min\{k_1, n-k_1\}}\binom{n}{k_1, k_2, n-k_1-k_2}\alpha(A)^{k_1}\alpha(B)^{k_2}(1-\alpha(A)-\alpha(B))^{n-k_1-k_2}\\
p(A|ABT_B) =& \frac{1}{2}\sum_{k=0}^{\lfloor\frac{n-1}{2}\rfloor}\binom{n}{k, k+1, n-2k-1}\beta(B)^k\beta(A)^{k+1}(1-\beta(A)-\beta(B))^{n-2k-1}\\
&+\sum_{k_1=2}^{n}\sum_{k_2=0}^{\min\{k_1-2, n-k_1\}}\binom{n}{k_1, k_2, n-k_1-k_2}\beta(A)^{k_1}\beta(B)^{k_2}(1-\beta(A)-\beta(B))^{n-k_1-k_2}
\end{aligned}$$

Denote $\underline{p}$ and $\bar{p}$ as a lower bound for $p(A|ABT_A)$ and an upper bound for $p(A|ABT_B)$ respectively, hence $\underline{p} > \bar{p} \implies p(A|ABT_A) > p(A|ABT_B)$, the attraction effect:

$$\underline{p} = \sum_{k_1=1}^{n}\sum_{k_2=0}^{\min\{k_1-1, n-k_1\}}\binom{n}{k_1, k_2, n-k_1-k_2}\alpha(A)^{k_1}\alpha(B)^{k_2}(1-\alpha(A)-\alpha(B))^{n-k_1-k_2}$$

$$\bar{p} = \sum_{k_1=1}^{n} \sum_{k_2=0}^{\min\{k_1-1, n-k_1\}} \binom{n}{k_1, k_2, n-k_1-k_2} \beta(A)^{k_1} \beta(B)^{k_2} (1-\beta(A)-\beta(B))^{n-k_1-k_2}$$

so $\underline{p} > \bar{p}$ if for $k_1 > k_2$:

$$\alpha(A)^{k_1}\alpha(B)^{k_2}(1-\alpha(A)-\alpha(B))^{n-k_1-k_2} > \beta(A)^{k_1}\beta(B)^{k_2}(1-\beta(A)-\beta(B))^{n-k_1-k_2}$$

and if $\mathcal{A}(A) = \mathcal{B}(B)$, $\mathcal{A}(B) = \mathcal{B}(A)$ and $\mathcal{A}(A, B) = \mathcal{B}(A, B)$ this simplifies to

$$\left[\frac{\alpha(A)}{\alpha(B)}\right]^{k_1-k_2} \left[= \left[\frac{\beta(B)}{\beta(A)}\right]^{k_1-k_2}\right] > 1. \tag{B1}$$

Due to the symmetry conditions, repeating the procedure for $p(B|ABT_B)$ and $p(B|ABT_A)$ gives the same expression. When in addition one lets $\frac{\alpha(A)}{\alpha(B)} \left[= \frac{\beta(B)}{\beta(A)}\right] > 1$, (B1) is satisfied as $k_1 > k_2$. ■

Proof of Proposition 2:

$\alpha(A) = \mathcal{A}(A) + \frac{1}{2}\left(\mathcal{A}(A, B) + \mathcal{A}(A, C_A)\right) + \frac{1}{3}\mathcal{A}(A, B, C_A)$

$\alpha(B) = \mathcal{A}(B) + \frac{1}{2}\mathcal{A}(A, B) + \frac{1}{3}\mathcal{A}(A, B, C_A)$

$\alpha(C_A) = \mathcal{A}(C_A) + \frac{1}{2}\mathcal{A}(A, C_A) + \frac{1}{3}\mathcal{A}(A, B, C_A)$

$\alpha_0 = 1 - \alpha(A) - \alpha(B) - \alpha(C_A)$

$\beta(A) = \mathcal{B}(A) + \frac{1}{2}\mathcal{B}(A, B) + \frac{1}{3}\mathcal{B}(A, B, C_B)$

$\beta(B) = \mathcal{B}(B) + \frac{1}{2}\left(\mathcal{B}(A, B) + \mathcal{B}(B, C_B)\right) + \frac{1}{3}\mathcal{B}(A, B, C_B)$

$\beta(C_B) = \mathcal{B}(C_B) + \frac{1}{2}\mathcal{B}(B, C_B) + \frac{1}{3}\mathcal{B}(A, B, C_B)$

$\beta_0 = 1 - \beta(A) - \beta(B) - \beta(C_B)$

These allow us to write:

$$p(A|ABC_A) = \frac{1}{3}\sum_{k=0}^{\lfloor \frac{n}{3} \rfloor} \binom{n}{k, k, k, n-3k} \alpha(A)^k \alpha(B)^k \alpha(C_A)^k \alpha_0^{n-3k}$$
$$+ \frac{1}{2}\sum_{k_1=1}^{\lfloor \frac{n}{2} \rfloor} \sum_{k_2=0}^{\min\{k_1-1, n-2k_1-1\}} \binom{n}{k_1, k_1, k_2, n-2k_1-k_2} \alpha(A)^{k_1} \alpha(B)^{k_1} \alpha(C_A)^{k_2} \alpha_0^{n-2k_1-k_2}$$

$$+\frac{1}{2}\sum_{k_1=1}^{\lfloor \frac{n}{2} \rfloor}\sum_{k_2=0}^{\min\{k_1-1,n-2k_1-1\}}\binom{n}{k_1,k_2,k_1,n-2k_1-k_2}\alpha(A)^{k_1}\alpha(B)^{k_2}\alpha(C_A)^{k_1}\alpha_0^{n-2k_1-k_2}$$
$$+\sum_{k_1=1}^{n}\sum_{k_2=0}^{\min\{k_1-1,n-k_1\}}\sum_{k_3=0}^{\min\{k_1-1,n-k_1-k_2\}}\binom{n}{k_1,k_2,k_3,n-k_1-k_2-k_3}\alpha(A)^{k_1}\alpha(B)^{k_2}\alpha(C_A)^{k_3}\alpha_0^{n-k_1-k_2-k_3}$$

$$p(A|ABC_B)=\frac{1}{3}\sum_{k=0}^{\lfloor \frac{n}{3} \rfloor}\binom{n}{k,k,k,n-3k}\beta(A)^{k}\beta(B)^{k}\beta(C_B)^{k}\beta_0^{n-3k}$$
$$+\frac{1}{2}\sum_{k_1=1}^{\lfloor \frac{n}{2} \rfloor}\sum_{k_2=0}^{\min\{k_1-1,n-2k_1-1\}}\binom{n}{k_1,k_1,k_2,n-2k_1-k_2}\beta(A)^{k_1}\beta(B)^{k_1}\beta(C_B)^{k_2}\beta_0^{n-2k_1-k_2}$$
$$+\frac{1}{2}\sum_{k_1=1}^{\lfloor \frac{n}{2} \rfloor}\sum_{k_2=0}^{\min\{k_1-1,n-2k_1-1\}}\binom{n}{k_1,k_2,k_1,n-2k_1-k_2}\beta(A)^{k_1}\beta(B)^{k_2}\beta(C_B)^{k_1}\beta_0^{n-2k_1-k_2}$$
$$+\sum_{k_1=1}^{n}\sum_{k_2=0}^{\min\{k_1-1,n-k_1\}}\sum_{k_3=0}^{\min\{k_1-1,n-k_1-k_2\}}\binom{n}{k_1,k_2,k_3,n-k_1-k_2-k_3}\beta(A)^{k_1}\beta(B)^{k_2}\beta(C_B)^{k_3}\beta_0^{n-k_1-k_2-k_3}$$

We make the following symmetry assumptions: $\mathcal{A}(A)=\mathcal{B}(B)$, $\mathcal{A}(B)=\mathcal{A}(C_A)=\mathcal{B}(A)=\mathcal{B}(C_B)$, $\mathcal{A}(A,B)=\mathcal{A}(A,C_A)=\mathcal{B}(A,B)=\mathcal{B}(B,C_B)$ and $\mathcal{A}(A,B,C_A)=\mathcal{B}(A,B,C_B)$. Together these imply $\alpha(A)=\beta(B)$, $\alpha(B)=\beta(A)$, $\alpha(C_A)=\beta(C_B)$ and $\alpha_0=\beta_0$. Under these assumptions, notice that $p(A|ABC_A)>p(A|ABC_B)$ (the compromise effect) is found if:

$$\left[\frac{\alpha(A)}{\alpha(B)}\right]^{k_1-k_2}\left[=\left[\frac{\beta(B)}{\beta(A)}\right]^{k_1-k_2}\right]>1. \tag{B2}$$

Due to the symmetry conditions, repeating the procedure for $p(B|ABC_B)$ and $p(B|ABC_A)$ gives the same expression. Finally, by symmetry,

$$\frac{\mathcal{A}(A)+\frac{1}{2}\mathcal{A}(A,B)+\frac{1}{2}\mathcal{A}(A,C_A)}{\mathcal{B}(A)+\frac{1}{2}\mathcal{B}(A,B)}>1 \implies \frac{\alpha(A)}{\alpha(B)}>1 \implies \text{(B2) holds as } k_1>k_2. \ \blacksquare$$

Proof of Proposition 3:

$\alpha(A)=\mathcal{A}(A)+\frac{1}{2}\mathcal{A}(A,S_A)+\frac{1}{3}\mathcal{A}(A,B,S_A)$

$\alpha(B)=\mathcal{A}(B)+\frac{1}{2}\mathcal{A}(B,S_A)+\frac{1}{3}\mathcal{A}(A,B,S_A)$

$\alpha(S_A)=\mathcal{A}(S_A)+\frac{1}{2}\left(\mathcal{A}(A,S_A)+\mathcal{A}(B,S_A)\right)+\frac{1}{3}\mathcal{A}(A,B,S_A)$

$$\alpha_0 = 1 - \alpha(A) - \alpha(B) - \alpha(S_A)$$

$$\beta(A) = \mathcal{B}(A) + \tfrac{1}{2}\mathcal{B}(A, S_B) + \tfrac{1}{3}\mathcal{B}(A, B, S_B)$$

$$\beta(B) = \mathcal{B}(B) + \tfrac{1}{2}\mathcal{B}(B, S_B) + \tfrac{1}{3}\mathcal{B}(A, B, S_B)$$

$$\beta(S_B) = \mathcal{B}(S_B) + \tfrac{1}{2}\left(\mathcal{B}(A, S_B) + \mathcal{B}(B, S_B)\right) + \tfrac{1}{3}\mathcal{B}(A, B, S_B)$$

$$\beta_0 = 1 - \beta(A) - \beta(B) - \beta(S_A)$$

These allow us to write:

$$\begin{aligned}
p(A|ABS_A) &= \frac{1}{3}\sum_{k=0}^{\lfloor \frac{n}{3} \rfloor}\binom{n}{k, k, k, n-3k}\alpha(A)^k\alpha(B)^k\alpha(S_A)^k\alpha_0^{n-3k} \\
&+ \frac{1}{2}\sum_{k_1=1}^{\lfloor \frac{n}{2} \rfloor}\sum_{k_2=0}^{\min\{k_1-1, n-2k_1-1\}}\binom{n}{k_1, k_1, k_2, n-2k_1-k_2}\alpha(A)^{k_1}\alpha(B)^{k_1}\alpha(S_A)^{k_2}\alpha_0^{n-2k_1-k_2} \\
&+ \frac{1}{2}\sum_{k_1=1}^{\lfloor \frac{n}{2} \rfloor}\sum_{k_2=0}^{\min\{k_1-1, n-2k_1-1\}}\binom{n}{k_1, k_2, k_1, n-2k_1-k_2}\alpha(A)^{k_1}\alpha(B)^{k_2}\alpha(S_A)^{k_1}\alpha_0^{n-2k_1-k_2} \\
&+ \sum_{k_1=1}^{n}\sum_{k_2=0}^{\min\{k_1-1, n-k_1\}}\sum_{k_3=0}^{\min\{k_1-1, n-k_1-k_2\}}\binom{n}{k_1, k_2, k_3, n-k_1-k_2-k_3}\alpha(A)^{k_1}\alpha(B)^{k_2}\alpha(S_A)^{k_3}\alpha_0^{n-k_1-k_2-k_3}
\end{aligned}$$

$$\begin{aligned}
p(A|ABS_B) &= \frac{1}{3}\sum_{k=0}^{\lfloor \frac{n}{3} \rfloor}\binom{n}{k, k, k, n-3k}\beta(A)^k\beta(B)^k\beta(S_B)^k\beta_0^{n-3k} \\
&+ \frac{1}{2}\sum_{k_1=1}^{\lfloor \frac{n}{2} \rfloor}\sum_{k_2=0}^{\min\{k_1-1, n-2k_1-1\}}\binom{n}{k_1, k_1, k_2, n-2k_1-k_2}\beta(A)^{k_1}\beta(B)^{k_1}\beta(S_B)^{k_2}\beta_0^{n-2k_1-k_2} \\
&+ \frac{1}{2}\sum_{k_1=1}^{\lfloor \frac{n}{2} \rfloor}\sum_{k_2=0}^{\min\{k_1-1, n-2k_1-1\}}\binom{n}{k_1, k_2, k_1, n-2k_1-k_2}\beta(A)^{k_1}\beta(B)^{k_2}\beta(S_B)^{k_1}\beta_0^{n-2k_1-k_2} \\
&+ \sum_{k_1=1}^{n}\sum_{k_2=0}^{\min\{k_1-1, n-k_1\}}\sum_{k_3=0}^{\min\{k_1-1, n-k_1-k_2\}}\binom{n}{k_1, k_2, k_3, n-k_1-k_2-k_3}\beta(A)^{k_1}\beta(B)^{k_2}\beta(S_B)^{k_3}\beta_0^{n-k_1-k_2-k_3}
\end{aligned}$$

We make the following symmetry assumptions: $\mathcal{A}(A) = \mathcal{B}(B)$, $\mathcal{A}(B) = \mathcal{B}(A)$, $\mathcal{A}(S_A) = \mathcal{B}(S_B)$, $\mathcal{A}(A, S_A) = \mathcal{B}(B, S_B)$, $\mathcal{A}(B, S_A) = \mathcal{B}(A, S_B)$ and $\mathcal{A}(A, B, S_A) = \mathcal{B}(A, B, S_B)$. Together these imply $\alpha(A) = \beta(B)$, $\alpha(B) = \beta(A)$, $\alpha(S_A) = \beta(S_B)$ and $\alpha_0 = \beta_0$. Under these assumptions, notice that $p(A|ABS_A) > p(A|ABS_B)$ (the similarity effect) is found if:

$$\left[\frac{\alpha(A)}{\alpha(B)}\right]^{k_1-k_2}\left[= \left[\frac{\beta(B)}{\beta(A)}\right]^{k_1-k_2}\right] > 1. \tag{B3}$$

Due to the symmetry conditions, repeating the procedure for $p(B|ABS_B)$ and $p(B|ABS_A)$ gives the same expression. Finally, by symmetry,

$$\frac{\mathcal{A}(A) + \frac{1}{2}\mathcal{A}(A, S_A)}{\mathcal{B}(A) + \frac{1}{2}\mathcal{B}(A, S_B)} > 1 \implies \frac{\alpha(A)}{\alpha(B)} > 1 \implies \text{(B3) holds as } k_1 > k_2. \ \blacksquare$$

References

Aldrovandi, S., Wood, A.M. and Brown, G.D.A. (2013). 'Sentencing, severity, and social norms: A rank-based model of contextual influence on judgments of crimes and punishments', *Acta Psychologica*, vol. 144(3), pp. 538–547.

Aldrovandi, S., Wood, A.M., Maltby, J. and Brown, G.D.A. (2014). 'Students' concern about indebtedness: a rank based social norms account', *Studies in Higher Education*, (ahead-of-print), pp. 1–21.

Bachi, B. and Spiegler, R. (2015). 'Buridanic competition', Working paper, Tel Aviv University, Tel Aviv, Israel.

Bhatia, S. (2013). 'Associations and the accumulation of preference.', *Psychological Review*, vol. 120(3), p. 522.

Bordalo, P., Gennaioli, N. and Shleifer, A. (2012). 'Salience theory of choice under risk', *Quarterly Journal of Economics*, vol. 127(3), pp. 1243–1285.

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

Boyce, C.J., Brown, G.D.A. and Moore, S.C. (2010). 'Money and happiness rank of income, not income, affects life satisfaction', *Psychological Science*, vol. 21(4), pp. 471–475.

Brown, G.D.A., Gardner, J., Oswald, A.J. and Qian, J. (2008). 'Does wage rank affect employees' well-being?', *Industrial Relations: A Journal of Economy and Society*, vol. 47(3), pp. 355–389.

Brown, G.D.A., Sandborn, A.N., Aldrovandi, S. and Wood, A.M. (2015a). 'The stability of unstable preferences and judgments: A rank-based approach to anchoring.', *Unpublished*.

Brown, G.D.A., Wood, A.M., Ogden, R.S. and Maltby, J. (2015b). 'Do student evaluations of university reflect inaccurate beliefs or actual experience? A relative rank model', *Journal of Behavioral Decision Making*, vol. 28(1), pp. 14–26.

Chandler, J., Mueller, P. and Paolacci, G. (2014). 'Nonnaïveté among Amazon Mechanical Turk workers: Consequences and solutions for behavioral researchers', *Behavior Research Methods*, vol. 46(1), pp. 112–130.

Cowan, N. (2001). 'The magical number 4 in short-term memory: A reconsideration of mental storage capacity.', *Behavioral and Brain Sciences*, vol. 24(1), pp. 87–185.

Doyle, J.R., O'Connor, D.J., Reynolds, G.M. and Bottomley, P.A. (1999). 'The robustness of the asymmetrically dominated effect: Buying frames, phantom alternatives, and in-store purchases', *Psychology & Marketing*, vol. 16(3), pp. 225–243.

Dubuisson, M.P. and Jain, A.K. (1994). 'A modified Hausdorff distance for object matching', in (*12th International Conference on Pattern Recognition*pp. 566–568, vol. 1, IEEE.

Fasano, G. and Franceschini, A. (1987). 'A multidimensional version of the Kolmogorov–Smirnov test', *Monthly Notices of the Royal Astronomical Society*, vol. 225(1), pp. 155–170.

Fiedler, K. (2000). 'Beware of samples! A cognitive-ecological sampling approach to judgment biases.', *Psychological Review*, vol. 107(4), p. 659.

K. Fiedler and P. Juslin, eds. (2006). *Information sampling and adaptive cognition*, Cambridge University Press.

Fiedler, K. and Kareev, Y. (2006). 'Does decision quality (always) increase with the size of information samples? Some vicissitudes in applying the law of large numbers.', *Journal of Experimental Psychology: Learning, Memory, and Cognition*, vol. 32(4), p. 883.

Frederick, S., Lee, L. and Baskin, E. (2014). 'The limits of attraction', *Journal of Marketing Research*, vol. 51(4), pp. 487–507.

Gennaioli, N. and Shleifer, A. (2010). 'What comes to mind', *Quarterly Journal of Economics*, vol. 125(4), pp. 1399–1433.

Gershoff, A.D. and Burson, K.A. (2011). 'Knowing where they stand: The role of inferred distributions of others in misestimates of relative standing', *Journal of Consumer Research*, vol. 38(3), pp. 407–419.

Goodman, J.K., Cryder, C.E. and Cheema, A. (2013). 'Data collection in a flat world: The strengths and weaknesses of Mechanical Turk samples', *Journal of Behavioral Decision Making*, vol. 26(3), pp. 213–224.

Horton, J.J., Rand, D.G. and Zeckhauser, R.J. (2011). 'The online laboratory: Conducting experiments in a real labor market', *Experimental Economics*, vol. 14(3), pp. 399–425.

Huber, J., Payne, J.W. and Puto, C. (1982). 'Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis', *Journal of Consumer Research*, vol. 9(1), pp. 90–98.

Huber, J., Payne, J.W. and Puto, C.P. (2014). 'Let's be honest about the attraction effect', *Journal of Marketing Research*, vol. 51(4), pp. 520–525.

Kamenica, E. (2008). 'Contextual inference in markets: On the informational content of product lines', *The American Economic Review*, vol. 98(5), pp. 2127–2149.

Kornienko, T. (2013). 'Nature's measuring tape: A cognitive basis for adaptive utility.', Working paper, University of Edinburgh, Edinburgh, UK.

Lindskog, M., Winman, A. and Juslin, P. (2013). 'Naïve point estimation.', *Journal of Experimental Psychology: Learning, Memory, and Cognition*, vol. 39(3), p. 782.

Maltby, J., Wood, A.M., Vlaev, I., Taylor, M.J. and Brown, G.D.A. (2012). 'Contextual effects on the perceived health benefits of exercise: The exercise rank hypothesis', *Journal of Sport & Exercise Psychology*, vol. 34(6), pp. 828–841.

Marsh, H.W. and Parducci, A. (1978). 'Natural anchoring at the neutral point of category rating scales', *Journal of Experimental Social Psychology*, vol. 14(2), pp. 193–204.

Maylor, E.A. and Roberts, M.A.J. (2007). 'Similarity and attraction effects in episodic memory judgments', *Cognition*, vol. 105(3), pp. 715–723.

Mellers, B.A. (1982). 'Equity judgment: A revision of Aristotelian views.', *Journal of Experimental Psychology: General*, vol. 111(2), p. 242.

Melrose, K.L., Brown, G.D.A. and Wood, A.M. (2013). 'Am I abnormal? Relative rank and social norm effects in judgments of anxiety and depression symptom severity', *Journal of Behavioral Decision Making*, vol. 26(2), pp. 174–184.

Mullainathan, S. (2002). 'A memory-based model of bounded rationality', *Quarterly Journal of Economics*, vol. 117(3), pp. 735–774.

Niedrich, R.W., Sharma, S. and Wedell, D.H. (2001). 'Reference price and price perceptions: A comparison of alternative models', *Journal of Consumer Research*, vol. 28(3), pp. 339–354.

Niedrich, R.W., Weathers, D., Hill, R.C. and Bell, D.R. (2009). 'Specifying price judgments with range-frequency theory in models of brand choice', *Journal of Marketing Research*, vol. 46(5), pp. 693–702.

Noguchi, T. and Stewart, N. (2014). 'In the attraction, compromise, and similarity effects, alternatives are repeatedly compared in pairs on single dimensions', *Cognition*, vol. 132(1), pp. 44–56.

Olivola, C.Y. and Sagara, N. (2009). 'Distributions of observed death tolls govern sensitivity to human fatalities', *Proceedings of the National Academy of Sciences*, vol. 106(52), pp. 22151–22156.

Paolacci, G. and Chandler, J. (2014). 'Inside the turk: Understanding Mechanical Turk as a participant pool', *Current Directions in Psychological Science*, vol. 23(3), pp. 184–188.

Paolacci, G., Chandler, J. and Ipeirotis, P.G. (2010). 'Running experiments on Amazon Mechanical Turk', *Judgment and Decision making*, vol. 5(5), pp. 411–419.

Papi, M. (2014). 'Noncompensatory consideration and compensatory choice: An application to Stackelberg competition', *Economic Theory Bulletin*, vol. 2(1), pp. 53–63.

Parducci, A., Calfee, R.C., Marshall, L.M. and Davidson, L.P. (1960). 'Context effects in judgment: Adaptation level as a function of the mean, midpoint, and median of the stimuli.', *Journal of Experimental Psychology*, vol. 60(2), pp. 65–77.

Parducci, A. and Perrett, L.F. (1971). 'Category rating scales: Effects of relative spacing and frequency of stimulus values.', *Journal of Experimental Psychology*, vol. 89(2), pp. 427–452.

Pettibone, J.C. and Wedell, D.H. (2000). 'Examining models of nondominated decoy effects across judgment and choice', *Organizational Behavior and Human Decision Processes*, vol. 81(2), pp. 300–328.

Prelec, D., Wernerfelt, B. and Zettelmeyer, F. (1997). 'The role of inference in context effects: Inferring what you want from what is available', *Journal of Consumer Research*, vol. 24(1), pp. 118–126.

Riskey, D.R., Parducci, A. and Beauchamp, G.K. (1979). 'Effects of context in judgments of sweetness and pleasantness', *Perception & Psychophysics*, vol. 26(3), pp. 171–176.

Robson, A.J. (2001). ‘Why would nature give individuals utility functions?’, *Journal of Political Economy*, vol. 109(4), pp. 900–914.

Roe, R.M., Busemeyer, J.R. and Townsend, J.T. (2001). ‘Multialternative decision field theory: A dynamic connectionst model of decision making.’, *Psychological Review*, vol. 108(2), pp. 370–392.

Shenoy, P. and Yu, A.J. (2013). ‘Rational preference shifts in multi-attribute choice: What is fair?’, *Proceedings of the Cognitive Science Society*.

Sher, S. and McKenzie, C.R.M. (2014). ‘Options as information: Rational reversals of evaluation and preference.’, *Journal of Experimental Psychology: General*, vol. 143(3), pp. 1127–1143.

Simonson, I. (1989). ‘Choice based on reasons: The case of attraction and compromise effects’, *Journal of Consumer Research*, vol. 16(2), pp. 158–174.

Simonson, I. (2008). ‘Will I like a ‘medium’ pillow? Another look at constructed and inherent preferences’, *Journal of Consumer Psychology*, vol. 18(3), pp. 155–169.

Simonson, I. (2014). ‘Vices and virtues of misguided replications: The case of asymmetric dominance’, *Journal of Marketing Research*, vol. 51(4), pp. 514–519.

Simonson, I. and Tversky, A. (1992). ‘Choice in context: Tradeoff contrast and extremeness aversion.’, *Journal of Marketing Research*, vol. 29(3), pp. 281–295.

Smith, R.H., Diener, E. and Wedell, D.H. (1989). ‘Intrapersonal and social comparison determinants of happiness: A range-frequency analysis.’, *Journal of Personality and Social Psychology*, vol. 56(3), pp. 317–325.

Steiner, J. and Stewart, C. (2014). ‘Perceiving prospects properly.’, CEPR Discussion Papers No. 10123.

Stewart, N., Chater, N. and Brown, G.D.A. (2006). ‘Decision by sampling’, *Cognitive Psychology*, vol. 53(1), pp. 1–26.

Stewart, N., Reimers, S. and Harris, A.J.L. (2015). 'On the origin of utility, weighting, and discounting functions: How they get their shapes and how to change their shapes', *Management Science*, vol. 61(3), pp. 687–705.

Stewart, N. and Simpson, K. (2008). 'A decision-by-sampling account of decision under risk', in (N. Chater and M. Oaksford, eds.), *The probabilistic mind: Prospects for Bayesian cognitive science*, Oxford, England: Oxford University Press.

Trueblood, J.S. (2012). 'Multialternative context effects obtained using an inference task', *Psychonomic Bulletin & Review*, vol. 19(5), pp. 962–968.

Trueblood, J.S., Brown, S.D. and Heathcote, A. (2014). 'The multiattribute linear ballistic accumulator model of context effects in multialternative choice.', *Psychological Review*, vol. 121(2), pp. 179–205.

Trueblood, J.S., Brown, S.D., Heathcote, A. and Busemeyer, J.R. (2013). 'Not just for consumers: Context effects are fundamental to decision making', *Psychological Science*, vol. 24(6), pp. 901–908.

Tsetsos, K., Chater, N. and Usher, M. (2012). 'Salience driven value integration explains decision biases and preference reversal', *Proceedings of the National Academy of Sciences*, vol. 109(24), pp. 9659–9664.

Tversky, A. (1972). 'Elimination by aspects: A theory of choice', *Psychological Review*, vol. 79(4), pp. 281–299.

Ungemach, C., Stewart, N. and Reimers, S. (2011). 'How incidental values from our environment affect decisions about money, risk, and delay.', *Psychological Science*, vol. 22(2), pp. 253–260.

Usher, M. and McClelland, J.L. (2004). 'Loss aversion and inhibition in dynamical models of multialternative choice.', *Psychological Review*, vol. 111(3), pp. 757–769.

Vlaev, I., Chater, N., Stewart, N. and Brown, G.D.A. (2011). 'Does the brain calculate value?', *Trends in Cognitive Sciences*, vol. 15(11), pp. 546–554.

Wedell, D.H. (1991). 'Distinguishing among models of contextually induced preference reversals.', *Journal of Experimental Psychology: Learning, Memory, and Cognition*, vol. 17(4), pp. 767–778.

Wedell, D.H., Parducci, A. and Geiselman, R.E. (1987). 'A formal analysis of ratings of physical attractiveness: Successive contrast and simultaneous assimilation', *Journal of Experimental Social Psychology*, vol. 23(3), pp. 230–249.

Wedell, D.H., Santoyo, E.M. and Pettibone, J.C. (2005). 'The thick and the thin of it: Contextual effects in body perception', *Basic and Applied Social Psychology*, vol. 27(3), pp. 213–228.

Wernerfelt, B. (1995). 'A rational reconstruction of the compromise effect: Using market data to infer utilities', *Journal of Consumer Research*, vol. 21(4), pp. 627–633.

Wood, A.M., Brown, G.D.A., Maltby, J. and Watkinson, P. (2012). 'How are personality judgments made? A cognitive model of reference group effects, personality scale responses, and behavioral reactions', *Journal of Personality*, vol. 80(5), pp. 1275–1311.

Yang, S. and Lynn, M. (2014). 'More evidence challenging the robustness and usefulness of the attraction effect', *Journal of Marketing Research*, vol. 51(4), pp. 508–513.

E-Cigarettes: The Extent and Impact of Complementary Dual-Use

1. Introduction

As of 2014 there were 466 brands of e-cigarettes with sales of around $7 billion and rising (Evans, 2015). However, E-cigarettes remain highly controversial (Cork, 2009; Dreaper, 2014; Fairchild and Bayer, 2015; Tavernise, 2014; The Economist, 2014; Triggle, 2014). Protagonists stress their positive impact on increasing the smoking cessation rate whereas sceptics express concerns about how they might lead to the renormalisation of smoking, gateway effects and pose a possible health risk. In lieu of conclusive long-run studies, some medical surveys and smaller controlled trials have suggested e-cigarettes containing nicotine can be effective for smoking cessation, and can be more effective than conventional nicotine replacement therapies (e.g., Brown *et al.*, 2014a; see Rahman *et al.*, 2015 for a review). Academic research has not been unanimous on this point however, as shown in recent studies (Al-Delaimy *et al.*, 2015; Cressey, 2014, 2015). Many policy debates and studies have placed a strong emphasis on this benefit of e-cigarettes (Adriaens *et al.*, 2014; Cahn and Siegel, 2011; Hajek, 2014). However, many health professionals and organisations remain uncertain about the innovation e.g., the World Health Organization (2014) and the American College of Physicians (Crowley, 2015). One area identified by academics (Etter, 2015) and organisations e.g., the Food and Drug Administration (FDA) (2014a) and Centers for Disease Control and Prevention (CDC) (Clarke, 2015; Furlow, 2015) as particularly pressing for research is that of dual-use of electronic and traditional cigarettes.

Despite the controversy surrounding the industry, which is typically based on the uncertainty surrounding the relatively new product and the potential gateway effect on non-smokers (e.g., Dutra and Glantz, 2014), a vital question informed by an economic perspective remains largely unexplored: do e-cigarettes represent a true substitute for cigarette use or are they also perceived as a complementary product? The extent to which e-cigarettes are substitutes or complements will have significant implications for the design and execution of policy.

Instead of simply assuming that e-cigarettes are used by all dual-users as a substitute for regular cigarettes, it is important to consider that they may be complementary products for some; more specifically, dynamic complements (Berry *et al.*, 2014). Take for instance someone who smokes regular cigarettes but who is subject to regulation or social pressure. For example, perhaps they can no longer smoke in their workplace or in places where they meet and associate with friends or family. This increases the need to reduce consumption or to consume elsewhere, with the associated inconvenience. E-cigarettes may be unregulated and fulfil both the nicotine addiction associated with cigarette use and some of the social element. A cigarette user now has the option to smoke regular cigarettes where they can and complement this with the use of e-cigarettes where it is not possible (or appropriate) to smoke cigarettes.

Notwithstanding other arguments for and against their use, if e-cigarettes were only adopted as substitutes and increased the cessation rate of smokers, they could increase the associated health and financial benefits due to lowering regular cigarette consumption. But where e-cigarettes act as complements they could instead blunt regular anti-smoking regulation and prolong the use of conventional cigarettes.

To tackle these issues we conducted an online survey of US participants and found that 37% of smokers who use e-cigarettes view them primarily as a complementary product to traditional cigarettes, rather than a substitute. We also found that 54.8% of 'substitute types' described themselves as trying to quit, but only 39.6% of 'complement types'. We use these results along with publicly available US data to calibrate a new cost-benefit analysis of e-cigarettes. Among other results we estimate that complementarity can reduce the benefits due to e-cigarettes by as much as 57%, or $3.3-4.9bn p.a., relative to the case with no complementarity.

We also compare the responses of dual-users to those of non-smokers. While 37% of dual-users view them primarily as a complementary product, only 27% of non-smokers thought the complementarity motive would be stronger for e-cigarette users. This 'perception gap' indeed suggests that many would overestimate the benefits of e-cigarettes.

In an analysis of the characteristics of dual-users we show that 'complement types' are more likely to be using some other cessation method or product. This suggests that some dual-users may be trying to quit via some other method, while using e-cigarettes to complement their smoking habit. We also find that men are much more likely use e-cigarettes as complements than women.

The paper proceeds as follows: Section 2 presents a model to guide our empirical strategy; Section 3 describes the survey design, implementation and results; Section 4 presents the cost benefit analysis; Section 5 concludes. The Supplementary Appendix provides the survey transcript.

2. Model

As in any model, what we propose is of course a simplification of reality. Our approach is directed by an empirical strategy. There are economic and medical aspects of the smoking problem that are beyond the scope of this paper which we do not incorporate e.g., market structure, firm decisions, direct effects of e-cigarettes on health, etc.

We use a two-period setup to model the individual's choice problem. There is a unit mass of individuals who smoke. In the first period individuals choose whether or not to adopt e-cigarettes. By the second period, they will either have ceased smoking or not. If they have ceased, they receive a 'health benefit' of h. Expected 'consumption utility' is denoted u_k^i when only conventional cigarettes are consumed and u_d^i when both conventional and electronic cigarettes are consumed. Consumption utility does not just refer to the physical activity of inhalation, but is intended to have a broader definition to include the effects of regulation on cigarettes and e-cigarettes, price, etc.[1] We assume that smokers prefer the health benefits from quitting to continuing to be a smoker or dual-user:

[1]Regulation is an important and current issue for individuals considering adopting electronic cigarettes. The decision of whether to consume e-cigarettes cannot be made in isolation of the regulatory regime.

Assumption 1. *For each smoker i, the benefit of quitting is greater than the benefit of continuing to be a smoker or dual-user i.e.,* $h > u_k^i, u_d^i$.

Under this assumption, each individual prefers the outcome of quitting than that of continued smoking. Their decision is whether or not e-cigarette adoption offers the preferred way to achieve cessation.

Without e-cigarette use, individuals will attempt to quit with probability q_k, and conditional on quit-attempt, will successfully quit with probability γ_k. The cessation rate for smokers is denoted $Q_k = q_k\gamma_k$. Individuals face uncertainty about the impact of e-cigarette adoption. Specifically, each individual is unsure whether for them e-cigarettes will act as a complement or a substitute for conventional cigarettes. Each individual i has a probability p_c^i of becoming a complement type, and $1-p_c^i$ of becoming a substitute type. Complement and substitute types are permitted to have different probabilities of attempting to quit (denoted q_c and q_s), and different probabilities of successfully quitting conditional on an attempt, denoted γ_c and γ_s. The cessation rates for these types are denoted accordingly by $Q_c = q_c\gamma_c$ and $Q_s = q_s\gamma_s$. To interpret substitute types, we assume they have a higher chance of quitting:[2]

Assumption 2. *The quit rate for substitute types is highest i.e.,* $Q_s > Q_c, Q_k$.

This assumption is later confirmed by the data. In the first period's decision problem, the second period's payoff is weighted by the discount factor $\delta \in (0, 1)$. In the first period, smoker i adopts e-cigarettes if the expected utility from adopting is greater than not adopting:

$$u_d^i+\delta\left[p_c^i[Q_ch + (1 - Q_c)u_d^i] + (1 - p_c^i)[Q_sh + (1 - Q_s)u_d^i]\right] \geq u_k^i+\delta[Q_kh+(1-Q_k)u_k^i]$$

$$p_c^i \leq \bar{p}_c^i \equiv \frac{(1+\delta)(u_d^i - u_k^i) + \delta\left[Q_s(h - u_d^i) - Q_k(h - u_k^i)\right]}{\delta(Q_s - Q_c)(h - u_d^i)} \tag{1}$$

That is, smokers with $p_c^i \leq \bar{p}_c^i$ will adopt e-cigarettes and become dual-users, which determines the dual-use rate, d. The dual-use rate depends on many factors. Importantly for us, smokers with a higher p_c^i are ceteris paribus less likely to adopt e-cigarettes. Our

[2]The cessation rate is important ratio in health policy directed towards regulation of smoking and so it is essential that we place this at the centre of our modelling exercise.

first empirical task is to establish that there exist smokers with $p_c^i > 0$.[3] Furthermore, if society's goal is to minimise the number of smokers and if e-cigarettes offer a better chance of cessation, then understanding individual differences in p_c^i is important and can aid policy design. Our survey enables us to investigate these differences.

The Health Cost Savings of E-cigarettes

The cost savings to society generated by one smoker quitting is denoted B. Without e-cigarettes, the financial savings from smokers quitting is therefore:

$$BQ_k. \tag{2}$$

If complement types are assumed not to exist, the total financial benefit of e-cigarettes would be given by:

$$B\left[(1-d)Q_k + dQ_s\right] \tag{3}$$

where we consider the dual-use rate d to be given as the rate seen in practice. In the presence of complement types, this expression is instead:

$$B\left[(1-d)Q_k + d[p_cQ_c + (1-p_c)Q_s]\right] \tag{4}$$

where $p_c > 0$ denotes the proportion of complement types. The difference between (4) and (2) is the actual benefit of e-cigarettes when complement types are taken into account. The difference between (4) and (3) is the error in the estimate of the benefit of e-cigarettes when it is incorrectly assumed that $p_c = 0$. The ratio of the error by the actual benefit is given by:

$$p_c\left[\frac{Q_s - Q_c}{Q_s - Q_k}\right]. \tag{5}$$

This is the relative reduction in the expected cost savings of e-cigarettes when complementary use is accounted for. Under Assumption 2, (5) is positive when $p_c > 0$ and is larger when there are more complement types (p_c). It is also larger when complement types have a lower probability of cessation, Q_c. We use our survey results alongside

[3]Note $d > 0$ does not trivially imply that there exists i such that $p_c^i > 0$: $d > 0$ and $p_c^i = 0$ for all i is compatible with (1) because the RHS is negative for individuals where u_k^i is sufficiently large.

publicly available data to estimate (5) and consider various scenarios for the variables we cannot estimate. We also use estimates of the financial cost per smoker (B) and the number of smokers to estimate (5) in absolute (dollar) terms.

3. The Survey

Participants

While the model provides a useful framework, the key issues in this paper need to be addressed empirically. With this in mind we carried out a survey of attitudes towards smoking. 2,406 participants completed our survey. The survey took an average of 4 minutes 24 seconds to complete, for which respondents were compensated with \$0.50, an implied wage of \$6.82 per hour. 36 had not heard of e-cigarettes, and so were removed, this left 2,370 for analysis. They were recruited through a leading online survey platform, Amazon Mechanical Turk in April 2015.[4]

The Amazon Mechanical Turk Platform

The platform allows data from a diverse participant pool to be collected at relatively little cost, and as such has become a popular recruitment tool for social science researchers (Kuziemko *et al.*, 2015; Paolacci and Chandler, 2014). Compared to the general population, participants recruited through the Mechanical Turk (MTurk) platform tend to be younger, better educated and more likely to be female (Berinsky *et al.*, 2012; Ipeirotis, 2010; Paolacci *et al.*, 2010; Paolacci and Chandler, 2014). Despite some demographic differences, the quality of responses has been found to be as reliable as that found in controlled laboratory environments across a variety of domains (Berinsky *et al.*, 2012; Horton *et al.*, 2011). Our survey is simpler than economic or psychological experiments, and hence we believe even less likely to generate unreliable data.

Demographic information of the whole sample and of the subsample of dual-users is provided in Table 1. The age and income of our participants are in line with these

[4]The survey transcript is given in the Supplementary Appendix.

existing studies using Amazon's platform. We did however find a higher proportion of males (55%) than that typically estimated to be the US MTurk population parameter. This however is likely explained by the higher prevalence of smoking (CDC, 2011) and dual-usage (Etter and Eissenberg, 2015) among men who may have been attracted to the study due to the mention of smoking and e-cigarettes in the study description.

Table 1

Respondent Demographics

	All	Dual-users
N	2,406	413
Age, mean [s.d.]	34.4 [0.2]	33.2 [0.5]
18-25	546 (23)	94 (23)
26-30	595 (25)	112 (27)
31-40	667 (28)	118 (29)
41-50	317 (13)	60 (15)
51+	281 (12)	39 (9)
Gender		
Male	1,316 (55)	254 (61.5)
Female	1,090 (45)	159 (38.5)
Income \$k, mean	51.3	48.2

Notes. Data shown are frequencies (% by column), unless otherwise specified. 'Dual-users' are defined as those who reported: "Yes" to the questions "Have you smoked at least 100 cigarettes in your entire life?", "Do you now smoke cigarettes at all, no matter how regularly?" and "Do you now use e-cigarettes every day, some days, or not at all?" with either "Every day" or "Some days". Income is average household pre-tax income, where we assume midpoints are the average of each income band, and an average of \$150k for the >\$100k band.

Design

In the medical literature, some surveys have asked dual e-cigarette and cigarette users their motivation for using e-cigarettes (e.g., Adkison *et al.*, 2013; Brown *et al.*, 2014b; Dockrell *et al.*, 2013; Etter and Eissenberg, 2015; Goniewicz *et al.*, 2013; Li *et al.*, 2015; Tackett *et al.*, 2015; Zhu *et al.*, 2013). However, they simply provide a list of reasons and allow the respondent to answer each one "yes" or "no". To understand the relative importance of the reasons for e-cigarette use, it is necessary to require respondents to consider a trade-off between their reasons for using e-cigarettes. Here, a few studies have asked respondents to select the primary reason they started using e-cigarettes (Goniewicz *et al.*, 2013; Tackett *et al.*, 2015). We however, are the first to pose the question as a direct trade-off between two fundamental economic motivations and to provide information on the strength of this trade-off. To assess whether smokers who use e-cigarettes view them as a substitute or a complement, respondents were asked the question as shown in Figure 1.

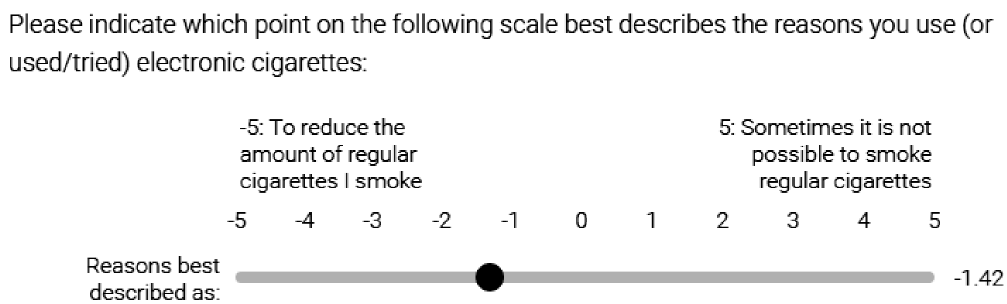

Figure 1. *Screen-shot of the Substitute or Complement Question*

Notes. This question was asked to those who reported ever having used an e-cigarette. As an example, this shot shows a participant selecting -1.42. Note that the cursor was centred at 0 when the page loaded to minimise bias and participants had to actively select a value before they could continue. There was no time limit for the question and participants could not go back to change their response once they had confirmed it.

The Extent of Complementary Dual-Use

413 of the 2,406 respondents reported being dual-users of electronic and regular cigarettes which we define as those who reported: "Yes" to the questions "Have you smoked at least 100 cigarettes in your entire life?", "Do you now smoke cigarettes at all, no matter how regularly?" and "Do you now use e-cigarettes every day, some days, or not at all?" with either "Every day" or "Some days". Responses to the question in Figure 1 are displayed in Figure 2.

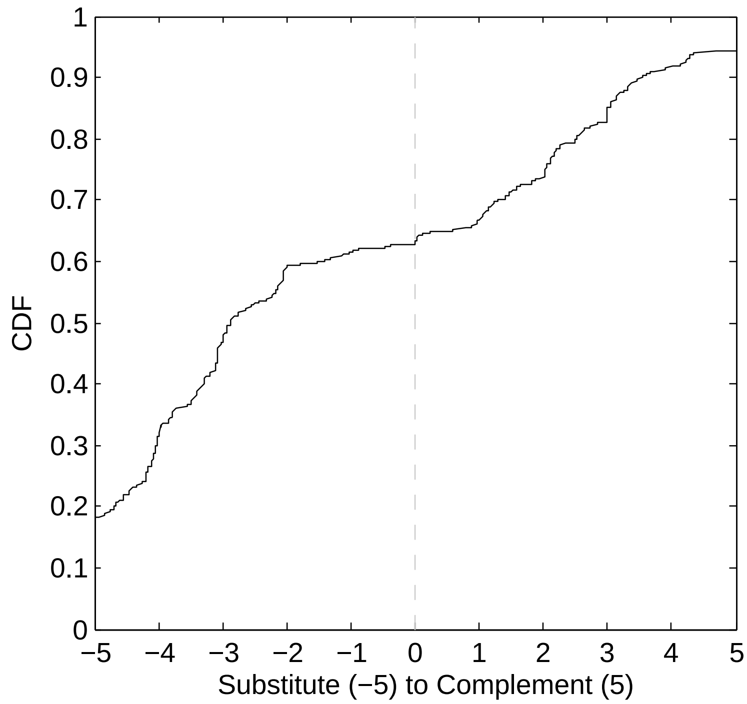

Figure 2. *Empirical CDF of Responses to the Question of Figure 1*

The data reveal that there is a large degree of variation in how users view these products. Figure 2 shows that 37% selected a point greater than 0 i.e., that they primarily use e-cigarettes as a complementary product. Moreover, 6% felt they were best described by the point furthest to the right, 5, suggesting that for these smokers, e-cigarettes are strong complements, dominating any substitution motive. There are then 63% who reported that they were best described by a point less than 0, indicating that they primarily use e-cigarettes to substitute away from traditional cigarette smoking. 18% felt they were best described by the point furthest to the left, -5, suggesting that for these smokers, e-cigarettes are strong substitutes. These data show that a majority of smokers use e-cigarettes primarily to reduce the amount they smoke. However, they also highlight that a sizable minority of e-cigarette users view them as complementary to traditional cigarettes.

As reported in the introduction, much public and media sentiment has emphasised the role of e-cigarettes as helping people quit smoking by acting as substitutes for conventional cigarettes. While our data support that many do indeed predominantly use e-cigarettes as substitutes, we have revealed that many in fact, do not. We now present evidence from our sample of this 'perception gap': that non-smokers believe there are significantly more substitute-type dual-users than there are (see Table 2). Non-smokers who had not tried e-cigarettes were asked the same question as in Figure 1 with the wording changed to: "Please indicate which point on the following scale best describes what you think the reasons are that people use electronic cigarettes". While 37.3% of dual-users reported using e-cigarettes primarily as complements to conventional cigarettes, only 26.8% of non-smokers thought the complementarity motive would be stronger for e-cigarette users. The difference of over 10% is highly significant [$\chi^2(1)$, p < 0.01]. Such a perception gap suggests that many would overestimate the benefits of e-cigarettes.

Table 2

The Perception Gap

Smoking status	Complement	Substitute
Non-smokers	26.8%	73.2%
Dual-users	37.3%	62.7%

Notes. There were 843 'non-smokers' who responded "no" to "Have you smoked at least 100 cigarettes in your entire life?", "Do you now smoke cigarettes at all, no matter how regularly?" and "Have you tried electronic cigarettes or "e-cigarettes", even just one time?". The 413 'dual-users' are as defined in the notes of Table 1. The difference between non-smokers and dual-users is significant: $\chi^2(1)$, p < 0.01.

We suggest that this perception gap may be due to the early success stories of those who ceased smoking. In Figure 3, we report a strong relationship between quit-status and why smokers use e-cigarettes. Ex-smokers who have successfully quit are the most likely to have used e-cigarettes primarily as substitutes (80%), followed by dual-users who are trying to quit (70%), and finally by those with no intention of quitting (56%).

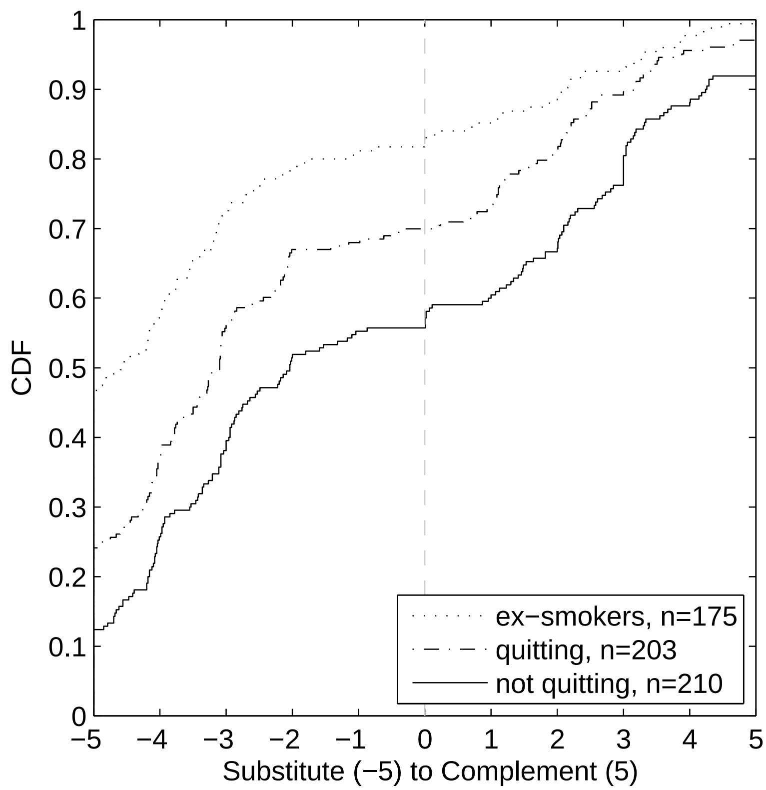

Figure 3. *Empirical CDF Decomposed by Quit Status*

Notes: 'ex-smokers' are those who reported having ceased smoking, but who have used e-cigarettes; 'quitting' are dual-users "currently trying to quit smoking cigarettes"; conversely 'not quitting' are dual-users who reported they were not. The distributions of responses display a clear ordering from left to right by quit-status: they are ordered by the statistical notion of first-order stochastic dominance. These differences are highly significant (p < 0.01 for all three pairwise KS tests).

That so many, and especially those who have no intention of quitting, use e-cigarettes primarily as a complementary product raises two concerns. Firstly, e-cigarettes may allow many smokers to continue their addiction to nicotine, and hence also to smoke regular cigarettes, more easily. Secondly, it suggests that public perception will not accurately reflect this if it is focused on the success stories of those who have used e-cigarettes and ceased smoking. It is likely that in the early years of e-cigarette usage the proportion of substitute users will appear relatively high, but as those desiring to quit do so, the proportion of smokers using e-cigarettes as a substitute in the long-run will necessarily be lower than it is now. This means that our finding of 63% from Figure 2 is likely an upper bound estimate on this proportion.

The Characteristics of Complement Types

We further investigate the differences in the characteristics of individuals who use e-cigarettes as a complement by the logistic regression results shown in Table 3. Confirming the results of Figure 3, we found that those reporting trying to quit are more likely to be those using e-cigarettes as a substitute (OR $= 0.45$, p < 0.01). We also found that if a smoker using e-cigarettes was also using another cessation method or product, it is much more likely that they use e-cigarettes as a complementary product (OR $= 2.7$, p < 0.01). This suggests that those who are trying to quit via some other product or method also use e-cigarettes, but do so to complement their smoking habit. Regarding levels of consumption, we found some evidence that those reporting higher e-cigarette usage are slightly more likely to be using them as a substitute (OR $= 0.97$, p $= 0.05$), whereas there was no difference in the quantity of cigarettes smoked. This does not necessarily mean that those reporting using e-cigarettes as a substitute are not successful in their attempts; it may well be that they are heavier smokers, and that e-cigarettes have helped them cut down. It is interesting however that any such substitution has only reduced their cigarette-consumption level to that of those using e-cigarettes as complements.

We found that being male is strongly associated with being a complement type (OR $= 6.5$, p < 0.01). This may relate to the stylised finding that men tend to engage in risky behaviour more than women (e.g., Harris *et al.*, 2006). There is no noteworthy difference across ages, but there is a small age effect by gender. Running separate regressions for men and women reveals no effect for men by age but that older women are slightly more likely to be complement types than younger women.[5] We found no relationship between income and the motivation for e-cigarette usage.

[5]These regressions are available from the authors.

Table 3

Characteristics of Complementary and Substitutive Dual-Users

	Log Odds (standard errors)	Odds Ratio [95% confidence interval]
Quitting	-0.866***	0.451***
	(0.231)	[0.267, 0.661]
Other Method(s)	1.007***	2.737***
	(0.281)	[1.578, 4.746]
E-cigarette Consumption	-0.027*	0.974*
	(0.014)	[0.948, 1.000]
Cigarette Consumption	0.003	1.003
	(0.014)	[0.976, 1.031]
Gender	1.874***	6.514***
	(0.701)	[1.649, 25.725]
Age	0.093	1.097
	(0.058)	[0.979, 1.229]
Age^2	-0.002*	0.998*
	(0.001)	[0.995, 1.000]
Gender*Age	-0.166**	0.847**
	(0.077)	[0.729, 0.985]
Gender*Age^2	0.003*	1.003*
	(0.002)	[1.000, 1.007]
Income	0.011	1.011
	(0.042)	[0.931, 1.097]
Constant	-1.237**	0.290**
	(0.616)	[0.087, 0.968]
N	413	
Pseudo R^2	0.12	

Notes. Dependent Variable $= 0$ if substitute, $= 1$ if complement. *Quitting*: response to "Currently trying to quit smoking cigarettes?" No $= 0$, Yes $= 1$; *Other Method(s)*: response to "Currently using another cessation service or product?" No $= 0$, Yes $= 1$; *E-cigarette Consumption*: response to "About how many times do you use e-cigarettes in a typical day?"; *Cigarette Consumption*: response to "About how many cigarettes do you smoke in a typical day?"; *Gender*: Female $= 0$, Male $= 1$; *Age:* was rescaled to start from zero by subtracting 18; *Income*: 10 ascending income bands; *Opinion*: about how e-cigarettes will affect society on average on a $[-5, 5]$ sliding scale to two decimal places with labels: $-5 =$ negative effect, $0 =$ no net effect, $5 =$ positive effect. Robust standard errors were used. ***p < 0.01, **p < 0.05, *p < 0.1.

4. Cost Benefit Analysis

We now assess how the opposing motivations of dual-use can affect the financial cost-benefit analysis of e-cigarettes. After outlining our procedure and findings, we provide a detailed walk-through of the methodology.

Financial cost-benefit analysis is a typical focus of economists and public health experts (World Health Organization, 2011; Chaloupka and Warner, 2000). Overwhelmingly, the evidence to date indicates that smoking has a net cost accounted for by health costs and reduced productivity (US Department of Health and Human Services, 2014b; Xu *et al.*, 2015). In the United States, 2013 official estimates imply that the annual net cost of smoking on the economy is approximately 2% of GDP, where each adult smoker imposes an annual cost on the economy of $6,865 (CDC, 2014; US Department of Health and Human Services, 2014b; Xu *et al.*, 2015). This represents the benefit per smoker quitting, B, from our model. If used purely for substitution, e-cigarettes could have a dramatically positive impact on the economy, which we estimate to be in the range of $5.7-8.5bn per annum.[6] However, if complementary use is ignored, then the number projected to cease smoking will be biased upwards and the expected cost-savings of e-cigarettes overestimated. We utilise our survey along with data on smoking habits to provide a quantification of the possible reduction in financial gains due to complementary use (see Table 4). We estimate the parameters of our model and estimate Equation (5). Doing so, we find that the presence of complementary use could reduce the financial benefit of e-cigarettes by up to 57% relative to the case where one assumes all e-cigarette users regard them as substitutes.

Our estimates are informed by our survey and available data. The 57% relative loss in benefits due to complementary use is our primary finding, and is independent of the number of dual-users (d) and the cost-savings per smoker quitting (B), as can be seen by inspection of (5). Rather, it depends on the quit-attempt rates and cease probabilities of those attempting to quit.

[6] Recall this excludes any potential health costs that might be directly attributable to e-cigarettes (e.g., Sussan *et al.*, 2015) or through the gateway effect (e.g., Dutra and Glantz, 2014).

In order to estimate the *absolute* benefit in monetary terms, we need an up-to-date estimate of the proportion of US smokers who are dual-users. Unfortunately recent officially-verified US data are not readily available. Our survey estimate (51%) is likely to be biased upwards because the survey's description specifically asked for opinions on smoking and e-cigarette use, attracting dual-users. A 2013 cross-country study (Adkison *et al.*, 2013) found that US former and current tobacco users rates of e-cigarette current-use, ever-use and awareness were roughly 1.5 times those in the UK. A 2015 study for Great Britain showed the dual-use rate among smokers to be 17.6% (Action on Smoking and Health, 2015). We therefore consider it reasonable to suppose that the US dual-use rate is higher than its UK counterpart, and if the same ratio has been preserved, to be around 26.4%. We suggest that the true current rate is likely to fall between these two, but could be even higher as e-cigarettes gain in popularity. In Table 4 we report both the more conservative estimate of 17.6% and the higher estimate of 26.4%, which yield reductions in benefits due to complementary use of up to \$3.3bn and \$4.9bn p.a. respectively.

Methodology

The economic impact assessment we conducted applies standard methodology taking account the costs of smoking and the benefits associated with cessation of smoking. The costs of smoking are extensively documented and were discussed in detail in the Surgeon General's 50th Anniversary Report (US Department of Health and Human Services, 2014b).

The CDC estimates that the cost of smoking imposes an annual net cost to the US economy in region of \$289-326 billion for 2013 (CDC, 2014; Xu *et al.*, 2015). To give conservative values, we utilise the lower bound of this range. The total number of adult smokers in the US is estimated at 17.8% of the adult population in 2013 (CDC, 2014). According to the CDC (2013) this translates into some 42.1 million smokers. These data allow us to derive a cost figure per adult smoker of \$6,865 per annum. In terms of our model, this is B. We assume that a person who ceases to smoke saves the economy this amount per annum. Hence, e-cigarettes, if they elevate the probability of cessation, present benefits. However, there are also costs associated with e-cigarettes, which include, among other things, a possible negative impact upon the probability of a smoker desiring to quit the habit. This effect is particularly more likely to be acute among those who treat e-cigarettes as complements to real cigarettes.

The impact e-cigarettes have upon the probability of cessation has not yet been quantified for the US. However, some, particularly those in the e-cigarette industry, suggest e-cigarettes are beneficial in this regard, a typical example is provided by the Quit Smoking Community (2015). By contrast, public health professionals in the US have generally been cautious with regard to such claims and some recent survey evidence indicates that e-cigarettes make quitting cigarette smoking less likely; which we agree with in respect of those who regard e-cigarettes as a complement. A recent study from the American Journal of Public Health (Al-Delaimy *et al.*, 2015) concludes:

> "Smokers who have used e-cigarettes may be at increased risk for not being able to quit smoking. These findings, which need to be confirmed by longer-term cohort studies, have important policy and regulation implications regarding the use of e-cigarettes among smokers."

Our online survey seeks to identify and distinguish between the two types of e-cigarette user: Substitute types who perceives e-cigarettes as a helpful aid to cessation and complement types who are likely to treat e-cigarettes as a consumption vehicle that obviates measures aimed at real cigarettes and thus allows for nicotine dependence to prevail. In determining how e-cigarettes may impact the probability of cessation we appealed to detailed survey data collected in England. The Smoking Toolkit Study (STS) by Smoking in England is a monthly survey of a nationally representative sample of adults that is designed to provide information about smoking prevalence and behaviour (also see Brown and West, 2015). Over the last five years or so the study has taken account of e-cigarettes. The STS is a monthly household survey that has been running since November 2006 and has accumulated more than 115,000 respondents. Each month about 1,800 are questioned and of these around 450 on average are smokers.

Data on e-cigarettes has been collected since the second quarter of 2011.

We analysed data collected in STS surveys and calculated the effect e-cigarettes have upon the probability of smoking cessation. This was done by making use of STS reported cessation probabilities. Cessation is a state where an ever-smoker has not smoked for at least one year. In England in 2011, the cessation rate of smokers was 4.4%, where the cessation probability conditional on attempting to quit in 2011 was 0.116. We assume these figures reflect the position before e-cigarettes became widely consumed: The prevalence of e-cigarettes at that time among smokers and recent ex-smokers was less than 3%. Data for 2014 gives the cessation rate at 0.066 and the probability of cessation conditional on attempting to quit at 0.175. E-cigarette prevalence was also higher, at 20.1%. The increase to 0.175 refers to all smokers attempting to quit. We assume that smokers in England who do not use e-cigarettes, continue as in the past with the probability 0.116. Given that 29.4% were using e-cigarettes to quit, one can show that dual-users of cigarettes and e-cigarettes who are attempting to quit have a 0.317 probability of cessation.

The higher cessation rates in England in 2014 than in 2011 are assumed to be entirely due to e-cigarettes in our analysis. Of course other variables affect cessation, including the price of cigarettes, which has increased over this period above the rate of inflation by a considerable margin, while GDP per capita has hardly changed (Office of National Statistics, 2015; The Statistics Portal, 2015). In addition, policy measures deterring smoking have changed and generally become more stringent. Legislation introduced in England in July 2007 banning smoking in public places has been strengthened in many instances by private initiatives since that date (for details see Smoking in England, 2015). We therefore likely overstate the impact of e-cigarettes. Our analysis does not claim to offer an exact absolute value for the effect of e-cigarettes, but rather demonstrates the possible differences between scenarios.

In 2010 the cessation rate in the US was 0.062 and the probability of cessation conditional on attempting to quit can be estimated at 0.118 (CDC, 2011). (In terms of our model, this serves as γ_k.) We employ this estimate for the probability of cessation conditional on attempting to quit for smokers who do not adopt e-cigarettes. We are unaware of data for the US that would allow us to compute the analogous figure to 0.317 as derived from English data, so we use the English estimate to inform our US scenarios, as explained below. We note that the estimated probabilities 0.116 and 0.118 are similar, as are smoking prevalence rates; as of 2014 these were 18.5% and 17.8% respectively (Smoking in England, 2015; CDC, 2014). Furthermore, both countries exhibit similar anti-smoking measures.

To compute the financial benefit to society of e-cigarettes, we assume that each smoker ceasing to smoke saves the US economy \$6,865 per annum (B). Pre e-cigarettes, we use the estimated cessation probability conditional on attempting to quit, $\gamma_k = 0.118$, along with the quit-attempt rate $q_k = 0.524$ (CDC, 2011), which gives the actual quit rate as $Q_k = 0.118 \times 0.524$. The results are shown in row 1 (and 6) of Table 4.

The first post e-cigarette scenario estimates the savings due to e-cigarettes where it is assumed that all dual-users are in fact substitute (S) types. Our survey found that 51% of smokers are dual-users. However, as discussed earlier, our survey is likely to have been answered by a high proportion of dual-users, so we consider a dual-use rate between 17.6-26.4% more plausible. Our survey also found that S-types have a quit-attempt rate of $q_s = 0.548$, which we do not consider a biased estimate. Conditional on attempting to quit, we assume that the cessation probability for S-types is $\gamma_s = 0.317$. We view this as a lower bound (and hence a conservative estimate) as it equals the gross rate we estimated for England, which itself will be a composite of complement (C) and S-type cessation probabilities. The results for this scenario are shown in row 2 (and 7) of Table 4. The difference in annual financial benefit between rows 1-2 (and 6-7) therefore reflects: (i) that e-cigarettes are associated with more smokers attempting to quit and (ii) that e-cigarettes give smokers attempting to quit an elevated chance of cessation. The benefit of e-cigarettes in this scenario is estimated to be \$5.7bn p.a. where $d = 0.176$, and \$8.5bn p.a. where $d = 0.264$, which we argue are likely to be substantially inflated if complementary use of e-cigarettes is ignored.

The latter post e-cigarette scenarios consider an increase to the cessation probability but distributed differently depending upon the type of dual-user. Our survey shows that 37% of dual-users are C-types, who have the lower quit-attempt rate of $q_c = 0.396$. We then assume that C-types are less likely to cease smoking than a substitute type given they attempt to quit. We take a parameter that scales this probability of cessation relative to the same probability for S-types, and for simplicity we report three values of it: 1, 0.5 and 0; corresponding to the values $\gamma_c = 0.317$, 0.158 and 0.[7] The results are reported in rows 3-5 (and 8-10) of Table 4. Using the extreme parameter value of 1, row 3 (and 8) produces a scenario where the cessation probability for complement and substitute types who are attempting to quit is identical. The difference in benefits between rows 2-3 of $0.9bn is then solely due to the lower rate at which complement types attempt to quit. Using the other extreme value of 0 in row 5 (and 10) posits that complement types have no chance of smoking cessation. Although extreme, it counterbalances the upward bias that attributes the cessation probability change over 2011-14 all to e-cigarettes, as discussed above. It is plausible then that the true value falls between these extremes, where we report the results corresponding to the parameter value of 0.5 in row 4 (and 9). The differences in benefits in rows 3-5 (and 8-10) are therefore solely due to the reduction in cessation probability of complement types. The central comparisons in benefits we make are those between the supposed benefit of e-cigarettes that many may assume, $5.7-8.5bn p.a., and the reduction in benefits implied by the presence of complement types i.e., the difference between row 2 (7) and one of rows 3,4,5 (8,9,10). These comparisons are reported in the final column of Table 4 where we find that the presence of complement types could reduce the financial benefit of e-cigarettes by up to $3.3bn p.a. under the dual-use rate 0.176, and $4.9p.a. under 0.264.

[7] Notice that these terms satisfy Assumption 2 as $Q_s = 0.174$, $Q_k = 0.062$ and $Q_c \in [0, 0.125]$

Table 4

Cost-Benefit Analysis Scenarios

Scenario	Dual-use rate	Proportion of C-types	Quit-attempt rates (smokers, S-types, C-types)	Cessation probability of those attempting to quit (smokers, S-types, C-types)	Relative reduction in cost-savings due to C-types	Number of smokers ceasing	Annual cost-savings in USD	Absolute reduction in cost-savings due to C-types in USD
	d	p_c	(q_k, q_s, q_c)	$(\gamma_k, \gamma_s, \gamma_c)$	(5)	n	nB	
Conservative dual-use rate estimate								
Pre EC	n/a	n/a	(0.524, n/a , n/a)	(0.118, n/a , n/a)	n/a	2.60m	17.9bn	n/a
Post EC with only S-type dual-users	0.176	0	(0.524, 0.548, n/a)	(0.118, 0.317, n/a)	n/a	3.43m	23.6bn	n/a
Post EC with S and C-type dual-users	0.176	0.37	(0.524, 0.548, 0.396)	(0.118, 0.317, 0.317)	16%	3.30m	22.6bn	0.9bn
	0.176	0.37	(0.524, 0.548, 0.396)	(0.118, 0.317, 0.158)	37%	3.13m	21.5bn	2.1bn
	0.176	0.37	(0.524, 0.548, 0.396)	(0.118, 0.317, 0.000)	57%	2.96m	20.3bn	3.3bn
Higher dual-use rate estimate								
Pre EC	n/a	n/a	(0.524, n/a , n/a)	(0.118, n/a , n/a)	n/a	2.60m	17.9bn	n/a
Post EC with only S-type dual-users	0.264	0	(0.524, 0.548, n/a)	(0.118, 0.317, n/a)	n/a	3.84m	26.4bn	n/a
Post EC with S and C-type dual-users	0.264	0.37	(0.524, 0.548, 0.396)	(0.118, 0.317, 0.317)	16%	3.65m	25.0bn	1.4bn
	0.264	0.37	(0.524, 0.548, 0.396)	(0.118, 0.317, 0.158)	37%	3.39m	23.3bn	3.1bn
	0.264	0.37	(0.524, 0.548, 0.396)	(0.118, 0.317, 0.000)	57%	3.13m	21.5bn	4.9bn

Notes. The terms in the second header row are from those from the model in Section 2.
Row 1 - Pre EC (E-cigarettes): Of smokers (no-EC) attempting to quit, 11.8% (2.60m) are successful, saving 2.60m $\times$ \$6,865 p.a.
Row 2 - Post EC with only substitute (S) types: All dual-users are S-types who have a quit-attempt rate of 54.8% as informed by our survey. Those attempting to quit are assumed to have the cessation probability 0.317.
Rows 3-5 - Post EC with S and complement (C) types: 37% of dual-users are C-types who have a quit-attempt rate of 39.6%, as informed by our survey. C-types attempting to quit are assumed to have either the same cessation probability than S-types (0.317), half (0.158), or a zero chance of cessation.

5. Conclusion

As detailed in the introduction, a great deal is made of the ability of e-cigarettes to help smokers to quit. To the extent that they are a viable substitute for conventional cigarettes this is undoubtedly true. Indeed, our work indicates that to most smokers they are a substitute. Combining this with the reportedly lower health risks entailed in smoking e-cigarettes, this generates a considerable benefit to society.

However, our study indicates that a substantial minority do not perceive e-cigarettes as a substitute but rather, as a complement. This can significantly dampen the benefits and may make some smokers worse off if they prolong smoking conventional cigarettes as a result, potentially undermining regulation. Consider for instance a smoker who is on the verge of quitting because of restrictions on smoking in the workplace, in social situations and perhaps because of the stigma attached to smoking. E-cigarettes allow such a smoker to continue smoking conventional cigarettes in private and in unregulated public spaces, but to use e-cigarettes to access oral nicotine delivery in situations, places or gatherings where real smoking is not possible.

We argue that longitudinal medical studies must not ignore the potential impact of complementarity. It is only through including concerns such as complementarity, alongside more visible concerns such as the potential gateway effect for non-smokers and any long-term health impact, that a full appraisal of the societal impact of e-cigarettes will be possible. In our own cost-benefit analysis we calculate that as much as 57% of the possible health benefits of e-cigarettes (absent concerns about gateway or uncertainties about the long-run health effects of e-cigarettes) may be removed through the effect of complementarity.

A special note of concern is the interplay between the role of e-cigarettes and the nature of regulation: increasing the degree of regulation of regular cigarettes will have a knock-on effect on e-cigarette use but also on how complementary they become. Ignoring this will result in a failure to predict correctly the impact of regulatory intervention.

Finally, we find that perceptions about complementarity differ significantly between non-smokers and dual-users. Dual-users are best placed to know the true level of complementarity since they have consumption experience of both e-cigarettes and conventional cigarettes. They return significantly higher levels of complementary use than non-smokers perceive. This difference in perceptions is especially worrying given that non-smokers make up the majority of the general population, likely including many policy-makers and health experts.

References

Action on Smoking and Health (2015). 'Use of electronic cigarettes (vapourisers) among adults in Great Britain', http://www.ash.org.uk/files/documents/ASH_891.pdf, (last accessed: 22 September 2015).

Adkison, S.E., O'Connor, R.J., Bansal-Travers, M., Hyland, A., Borland, R., Yong, H.H., Cummings, K.M., McNeill, A., Thrasher, J.F., Hammond, D. and Fong, G.T. (2013). 'Electronic nicotine delivery systems: International tobacco control four-country survey', *American Journal of Preventive Medicine*, vol. 44(3), pp. 207–215.

Adriaens, K., Van Gucht, D., Declerck, P. and Baeyens, F. (2014). 'Effectiveness of the electronic cigarette: An eight-week Flemish study with six-month follow-up on smoking reduction, craving and experienced benefits and complaints', *International Journal of Environmental Research and Public Health*, vol. 11(11), pp. 11220–11248.

Al-Delaimy, W.K., Myers, M.G., Leas, E.C., Strong, D.R. and Hofstetter, C.R. (2015). 'E-cigarette use in the past and quitting behavior in the future: A population-based study', *American Journal of Public Health*, vol. 105(6), pp. 1213–1219.

Berinsky, A.J., Huber, G.A. and Lenz, G.S. (2012). 'Evaluating online labor markets for experimental research: Amazon.com's Mechanical Turk', *Political Analysis*, vol. 20(3), pp. 351–368.

Berry, S., Khwaja, A., Kumar, V., Musalem, A., Wilbur, K.C., Allenby, G.M., Anand, B.N., Chintagunta, P.K., Hanemann, W.M., Jeziorski, P. and Mele, A. (2014). 'Structural models of complementary choices', *Marketing Letters*, vol. 25, pp. 245–256.

Brown, J., Beard, E., Kotz, D., Michie, S. and West, R. (2014a). 'Real-world effectiveness of e-cigarettes when used to aid smoking cessation: A cross-sectional population study', *Addiction*, vol. 109(9), pp. 1531–1540.

Brown, J. and West, R. (2015). 'Trends in electronic cigarette use in england', http://www.smokinginengland.info/latest-statistics, (last accessed: 1 July 2015).

Brown, J., West, R., Beard, E., Michie, S., Shahab, L. and McNeill, A. (2014b). 'Prevalence and characteristics of e-cigarette users in Great Britain: Findings from a general population survey of smokers', *Addictive Behaviors*, vol. 39(6), pp. 1120–1125.

Cahn, Z. and Siegel, M. (2011). 'Electronic cigarettes as a harm reduction strategy for tobacco control: A step forward or a repeat of past mistakes?', *Journal of Public Health Policy*, vol. 32(1), pp. 16–31.

Centers for Disease and Control and Prevention (2011). 'Morbidity and mortality weekly report', vol. 60(44), pp. 1513–1519.

Centers for Disease and Control and Prevention (2013). 'Current cigarette smoking among adults in the United States', http://www.cdc.gov/tobacco/data_statistics/fact_sheets/adult_data/cig_smoking/index.htm, (last accessed: 23 September 2015).

Centers for Disease and Control and Prevention (2014). 'Morbidity and mortality weekly report', vol. 63(47), pp. 1108–1112.

Chaloupka, F.J. and Warner, K.E. (2000). 'The economics of smoking', in (A. J. Culyer and J. P. Newhouse, eds.), *Handbook of health economics*, pp. 1539–1627, vol. 1B, chap. 29, Oxford, UK: Elsevier.

Clarke, T. (2015). 'U.S. ad campaign points to dangers of dual cigarette/e-cig use', *Reuters*.

Cork, K. (2009). ‘To vape or not to vape: Controversy swirls around e-cigarettes’, *National Association of Local Boards of Health*, vol. 16(4), p. 7.

Cressey, D. (2014). ‘Electronic cigarettes ‘don’t aid quitting’, study says’, *Nature News*, doi:10.1038/nature.2014.14918.

Cressey, D. (2015). ‘California’s anti-vaping bill goes up in smoke’, *Nature*, vol. 523(7560), p. 267.

Crowley, R.A. (2015). ‘Electronic nicotine delivery systems: Executive summary of a policy position paper from the American College of Physicians’, *Annals of Internal Medicine*, vol. 162(8), pp. 583–584.

Dockrell, M., Morison, R., Bauld, L. and McNeill, A. (2013). ‘E-cigarettes: Prevalence and attitudes in Great Britain’, *Nicotine & Tobacco Research*, vol. 15(10), pp. 1737–1744.

Dreaper, J. (2014). “Resist urge to control e-cigarettes’, WHO told’, *BBC News*.

Dutra, L.M. and Glantz, S.A. (2014). ‘Electronic cigarettes and conventional cigarette use among US adolescents: A cross-sectional study’, *JAMA Pediatrics*, vol. 168(7), pp. 610–617.

Etter, J.F. (2015). ‘E-cigarettes: Methodological and ideological issues and research priorities’, *BMC Medicine*, vol. 13(1), p. 32.

Etter, J.F. and Eissenberg, T. (2015). ‘Dependence levels in users of electronic cigarettes, nicotine gums and tobacco cigarettes’, *Drug and Alcohol Dependence*, vol. 147, pp. 68–75.

Evans, P. (2015). ‘E-cigarette makers face rise of counterfeits’, *The Wall Street Journal*.

Fairchild, A.L. and Bayer, R. (2015). ‘Smoke and fire over e-cigarettes’, *Science*, vol. 247(6220), pp. 375–376.

Furlow, B. (2015). 'US Government warns against long-term dual use of conventional and e-cigarettes', *The Lancet Respiratory Medicine*, vol. 3(5), p. 345.

Goniewicz, M.L., Lingas, E.O. and Hajek, P. (2013). 'Patterns of electronic cigarette use and user beliefs about their safety and benefits: An internet survey', *Drug and Alcohol Review*, vol. 32(2), pp. 133–140.

Hajek, P. (2014). 'Electronic cigarettes have a potential for huge public health benefit', *BMC Medicine*, vol. 12(1), p. 225.

Harris, C.R., Jenkins, M. and Glaser, D. (2006). 'Gender differences in risk assessment: Why do women take fewer risks than men?', *Judgment and Decision Making*, vol. 1(1), pp. 48–63.

Horton, J.J., Rand, D.G. and Zeckhauser, R.J. (2011). 'The online laboratory: Conducting experiments in a real labor market', *Experimental Economics*, vol. 14(3), pp. 399–425.

Ipeirotis, P.G. (2010). 'Demographics of Mechanical Turk', *CeDER working paper, New York University*.

Kuziemko, I., Norton, M.I., Saez, E. and Stantcheva, S. (2015). 'How elastic are preferences for redistribution? evidence from randomized survey experiments', *American Economic Review*, vol. 105(4), pp. 1478–1508.

Li, J., Newcombe, R. and Walton, D. (2015). 'The prevalence, correlates and reasons for using electronic cigarettes among New Zealand adults', *Addictive Behaviors*, vol. 45, pp. 245–251.

Office of National Statistics (2015). 'Statistical bulletin: Consumer price index', .

Paolacci, G. and Chandler, J. (2014). 'Inside the turk: Understanding Mechanical Turk as a participant pool', *Current Directions in Psychological Science*, vol. 23(3), pp. 184–188.

Paolacci, G., Chandler, J. and Ipeirotis, P.G. (2010). 'Running experiments on Amazon Mechanical Turk', *Judgment and Decision making*, vol. 5(5), pp. 411–419.

Quit Smoking Community (2015). 'Best electronic cigarette brands for 2015', http://www.quitsmokingcommunity.org/electronic-cigarettes, (last accessed: 23 September 2015).

Rahman, M.A., Hann, N., Wilson, A., Mnatzaganian, G. and Worrall-Carter, L. (2015). 'E-cigarettes and smoking cessation: Evidence from a systematic review and meta-analysis', *PLoS one*, vol. 10(3), doi:10.1371/journal.pone.0122544.

Smoking in England (2015). 'Smoking toolkit study', http://www.smokinginengland.info, (last accessed: 1 July 2015).

Sussan, T.E., Gajghate, S., Thimmulappa, R.K., Ma, J., Kim, J.H., Sudini, K., Consolini, N., Cormier, S.A., Lomnicki, S., Hasan, F., Pekosz, A. and Biswal, S. (2015). 'Exposure to electronic cigarettes impairs pulmonary anti-bacterial and anti-viral defenses in a mouse model', *PloS one*, vol. 10(2), doi:10.1371/journal.pone.0116861.

Tackett, A.P., Lechner, W.V., Meier, E., Grant, D.M., Driskill, L.M., Tahirkheli, N.N. and Wagener, T.L. (2015). 'Biochemically verified smoking cessation and vaping beliefs among vape store customers', *Addiction*, vol. 110(5), pp. 868–874.

Tavernise, S. (2014). 'Study gives e-cigarettes edge in helping smokers quit', *The New York Times*.

The Economist (2014). 'Why e-cigarettes are so controversial', http://www.economist.com/blogs/economist-explains/2014/05/economist-explains-2, (last accessed: 22 September 2015).

The Statistics Portal (2015). 'Consumer price index (cpi) of tobacco monthly in the United Kingdom (UK) from june 2013 to december 2014', http://www.statista.com/statistics/285027/tobacco-consumer-price-index-cpi-in-the-united-kingdom-uk/, (last accessed: 22 September 2015).

Triggle, N. (2014). ‘E-cigarettes ‘help smokers to quit”, *BBC News*.

US Department of Health and Human Services (2014a). *Federal Register*, vol. 79(80).

US Department of Health and Human Services (2014b). ‘The health consequences of smoking - 50 years of progress: A report of the surgeon general’, http://www.surgeongeneral.gov/library/reports/50-years-of-progress/, (last accessed: 22 September 2015).

World Health Organization (2011). ‘Economics of tobacco toolkit: Assessment of the economic costs of smoking’, http://apps.who.int/iris/bitstream/10665/44596/1/9789241501576_eng.pdf, (last accessed: 22 September 2015).

World Health Organization (2014). ‘Framework convention on tobacco control FCTC/COP/6/10’, http://apps.who.int/gb/fctc/PDF/cop6/FCTC_COP6_10-en.pdf?ua=1, (last accessed: 22 September 2015).

Xu, X., Bishop, E.E., Kennedy, S.M., Simpson, S.A. and Pechacek, T.F. (2015). ‘Annual healthcare spending attributable to cigarette smoking: An update’, *American Journal of Preventive Medicine*, vol. 48(3), pp. 326–333.

Zhu, S.H., Gamst, A., Lee, M., Cummins, S., Yin, L. and Zoref, L. (2013). ‘The use and perception of electronic cigarettes and snus among the US population’, *PLoS one*, vol. 8(10), doi:10.1371/journal.pone.0079332.

Supplementary Appendix

This appendix includes the transcript of the survey and a detailed description of the cost benefit analysis methodology. We would not anticipate either being in the final published paper but might be of use to referees and might be made available more generally in the form of an (online only) supplementary appendix.

The survey was conducted using the online survey software Qualtrics. The question numbers displayed here, along with the coded values shown in parentheses correspond to the coding of the data as it is displayed in the results file available online. The questions were organised in blocks. Which blocks participants saw depended on their prior answers. We indicate any criteria to see a block. Within blocks, some questions were also restricted to be shown only to participants giving particular prior answers. Where there are such conditions, they are shaded.

Block 1: All

Q23 In order to participate in this research study, it is necessary that you give your informed consent. By responding you are indicating that you understand the nature of the research study and your role and that you agree to participate in the research. Please consider the following points before continuing: I understand that I am participating in research conducted by the University of Warwick. I understand the research team will use anonymized data in any presentations of the research results. Data will not be associated with individuals and any identifying data will then be destroyed. I understand that my participation in this study is voluntary, and that after the study data collection has begun, I may refuse to participate further without any penalty. By continuing I am stating that I am over 18 years of age, and that I have read the above information and consent to participate in this study being conducted. Please click "I agree" to agree that you have read and understood the information above:

- ❍ I agree (1)

Block 2: All

Q24 Have you smoked at least 100 cigarettes in your entire life?

- ❍ Yes (1)
- ❍ No (2)

Q32 Do you now smoke cigarettes at all, no matter how regularly?

- ❍ Yes (1)
- ❍ No (2)

Block 3: If Q32 = Yes

Q25 Do you now smoke cigarettes every day or some days?

- ❍ Every day (1)
- ❍ Some days (2)

If Q25 = Every day:

Q27 About how many cigarettes do you smoke in a typical day?

If Q25 = Some days:

Q28 About how many cigarettes do you smoke in a typical week?

Block 4: If Q24=Yes and Q32=No

Q31 About how long has it been since you last smoked cigarettes?

- ❍ 0-3 months (1)
- ❍ 3-6 months (2)
- ❍ 6-12 months (3)
- ❍ 1-2 years (4)
- ❍ 2-3 years (5)
- ❍ 3-4 years (6)
- ❍ 4-5 years (7)

Q55 During the last period you smoked, did you smoke cigarettes every day or some days?
- ❍ Every day (1)
- ❍ Some days (2)

If Q55 = Every day

Q56 During this period, about how many cigarettes did you smoke in a typical day?

If Q55 = Some days Is Selected

Q57 During this period, about how many cigarettes do you smoke in a typical week?

Block 5: All

Q29 Have you tried Electronic Cigarettes or "E-cigarettes", even just one time?
- ❍ Yes (1)
- ❍ No (2)

If Q29 = No

Q40 Have you ever heard of Electronic Cigarettes or "E-cigarettes"?
- ❍ Yes (4)
- ❍ No (5)

Block 6: If Q29=Yes

Q30 Do you now use e-cigarettes every day, some days, or not at all?
- ❍ Every day (1)
- ❍ Some days (2)
- ❍ Not at all (3)

If Q30 = Every day

Q31 About how many times do you use e-cigarettes in a typical day?

If Q30 = Some days

Q32 About how many times do you use e-cigarettes in a typical week?

Block 7: If Q30=Not at all

Q33 About how long has it been since you last used e-cigarettes?
- ❍ 0-3 months (2)
- ❍ 3-6 months (3)
- ❍ 6-12 months (4)
- ❍ 1-2 years (5)
- ❍ 2-3 years (6)
- ❍ 3-4 years (7)
- ❍ 4-5 years (9)
- ❍ 5+ years (10)

Q65 During the last period you used e-cigarettes, did you use them every day, some days or just one time?
- ❍ Every day (1)
- ❍ Some days (2)
- ❍ Just one time (3)

If Q65 = Every day

Q66 During this period, about how many times do you use e-cigarettes in a typical day?

If Q65 = Some days

Q67 During this period, about how many times do you use e-cigarettes in a typical week?

Block 8: If Q29=Yes or Q40=Yes

Q41 No matter what your smoking history is, we are interested in your opinions.

Q11 Below are some of the advantages that people often think electronic cigarettes have over conventional cigarettes. Please give us your opinion of the order of importance of these advantages by ranking them from 1 (most important) to 6 (least important).

______ Lower health risks (9)
______ No second hand smoke (10)
______ Cheaper (11)
______ Can use them in many public places (14)
______ Less odor (12)
______ Less risk of causing a fire (13)

Q12 Below are some of the disadvantages that people often think electronic cigarettes have over conventional cigarettes. Please give us your opinion of the order of importance of these disadvantages by ranking them from 1 (most important) to 6 (least important).

______ Not the same experience as cigarettes (3)
______ Unsure about health risks (2)
______ Cost of equipment (1)
______ Too addictive (4)
______ Concern over product malfunction (5)
______ Confusing number of brands and products (6)

If Q29 = Yes

Q23 Please indicate which point on the following scale best describes the reasons you use (or used/tried) electronic cigarettes:

Reasons best described as: [-5,5] sliding scale to two decimal places as shown in Fig. 1. Label over -5 read “To reduce the amount of regular cigarettes I smoke”, label over 5 read “Sometimes it is not possible to smoke regular cigarettes”.

If Q29 = No

Q25 Please indicate which point on the following scale best describes what you think the reasons are that people use electronic cigarettes:
Reasons best described as: [-5,5] sliding scale to two decimal places. Label over -5 read “To reduce the amount of regular cigarettes they smoke”, label over 5 read “Sometimes it is not possible to smoke regular cigarettes”.

Q30 Considering the arguments for and against e-cigarettes, and that they may affect different people in different ways: Please use the slider below to indicate your opinion on how they will affect society on average?
My opinion is best described as: [-5,5] sliding scale to two decimal places. Label over -5 read “negative effect”, label over 0 read “no net effect”, label over 5 read “positive effect”.

Q31 If you would like to, please explain how you think e-cigarettes will impact society. (optional)

Block 9: Q32=Yes

Q32 Have you used any of the following smoking-cessation services or products? (select all that apply, if any)

	Smoking Cessation Methods	
	Currently (1)	In the past (2)
Doctor consultation (9)	❑	❑
Counseling (10)	❑	❑
Behavioral Therapy (11)	❑	❑
Other non-medical method (12)	❑	❑
Gum (2)	❑	❑
Patches (4)	❑	❑
Nasal Spray (5)	❑	❑
Inhalers (6)	❑	❑
Lozenges (7)	❑	❑
Other medication (8)	❑	❑

Q26 Would you describe yourself as someone who is currently trying to quit smoking cigarettes?

❍ Yes (2)

❍ No (3)

Block 10: If Q32=No and Q24=Yes

Q37 Have you used any of the following smoking-cessation services or products? (select all that apply, if any)

	Smoking Cessation Methods
	I have used the following: (1)
Doctor consultation (9)	❑
Counseling (10)	❑
Behavioral Therapy (11)	❑
Other non-medical method (12)	❑
Gum (2)	❑
Patches (4)	❑
Nasal Spray (5)	❑
Inhalers (6)	❑
Lozenges (7)	❑
Other medication (8)	❑

If Q29 = Yes and if Q37 has at least one item checked:

Q38 When you compare e-cigarettes to any other cessation methods you used, which was more effective for you?
Most effective for me: [-5,5] sliding scale to two decimal places. Label over -5 read “other methods”, label over 5 read “e-cigarettes”.

Block 11: All

Q43 Do you avoid certain unhealthy food or drink?
- Yes (1)
- No (2)

Q27 Are you currently a member of a health club or fitness center?
- Yes (1)
- No (2)

Q28 Do you pay into Medicare?
- Yes (1)
- No (2)
- Don't know (3)

Q29 Do you pay into a private health insurance plan?
- Yes (1)
- No (2)
- Don't Know (3)

Q15 Are you male or female?
- Male (1)
- Female (2)

Q16 How old are you?

Q31 Which state do you primarily live in?
- Alabama (1)
- Alaska (2)
- Arizona (3)
- Arkansas (4)
- California (5)
- Colorado (6)
- Connecticut (7)
- Delaware (8)
- District of Columbia (9)
- Florida (10)
- Georgia (11)
- Hawaii (12)
- Idaho (13)
- Illinois (14)
- Indiana (15)
- Iowa (16)
- Kansas (17)
- Kentucky (18)
- Louisiana (19)
- Maine (20)
- Maryland (21)
- Massachusetts (22)
- Michigan (23)
- Minnesota (24)
- Mississippi (25)
- Missouri (26)
- Montana (27)
- Nebraska (28)
- Nevada (29)
- New Hampshire (30)
- New Jersey (31)
- New Mexico (35)
- New York (32)
- North Carolina (33)
- North Dakota (34)
- Ohio (36)
- Oklahoma (37)
- Oregon (38)
- Pennsylvania (39)
- Rhode Island (40)
- South Carolina (41)
- South Dakota (42)
- Tennessee (43)

- Texas (44)
- Utah (45)
- Vermont (46)
- Virginia (47)
- Washington (48)
- West Virginia (49)
- Wisconsin (50)
- Wyoming (51)

Q17 Please indicate your household's annual pre-tax income. (US $)
- 0 - 9,999 (1)
- 10,000 - 19,999 (2)
- 20,000 - 29,999 (3)
- 30,000 - 39,999 (4)
- 40,000 - 49,999 (5)
- 50,000 - 59,999 (6)
- 60,000 - 69,999 (7)
- 70,000 - 79,999 (8)
- 80,000 - 99,999 (10)
- 100,000 or more (9)