

Original citation:

Proto, Eugenio and Sgroi, Daniel. (2017) Biased beliefs and imperfect information. Journal of Economic Behaviour and Organization, 136. pp. 186-202.

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/85608>

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions. Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

© 2017, Elsevier. Licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International <http://creativecommons.org/licenses/by-nc-nd/4.0/>

A note on versions:

The version presented here may differ from the published version or, version of record, if you wish to cite this item you are advised to consult the publisher's version. Please see the 'permanent WRAP URL' above for details on accessing the published version and note that access may require a subscription.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk

Biased Beliefs and Imperfect Information

Eugenio Proto ^a and Daniel Sgroi ^{b *}

^aUniversity of Warwick, CAGE and IZA ^bUniversity of Warwick and CAGE

January 2017

Abstract: We perform an experiment designed to assess the accuracy of beliefs about characteristics and decisions. Subjects are asked to declare beliefs typically formed through *real world experiences*. They are then asked to report beliefs concerning other individuals from the same environment. We test two main hypotheses: (i) whether for items not perfectly observable, individuals suffer from some type of biased beliefs; (ii) whether this bias is reduced when information is more readily available. We find a powerful and ubiquitous bias in perceptions that is “self-centered” in the sense that those at extremes tend to perceive themselves as closer to the middle of the distribution than is the case. This bias does not completely disappear when the information is more readily available. We present evidence from our experiment that *limited attention* and *self-serving deception* can provide explanations for this bias and present important economic applications.

Keywords: biased beliefs, information, attitudes, characteristics, self-centered bias

JEL classification: D03, C83, D84.

“Let them eat cake.” (Commonly attributed to) Marie Antoinette (1755–1793), Archduchess of Austria and Queen of France.¹

1 Introduction

Early work attempting to address the question of whether individuals hold biased beliefs linked information imperfection with potentially heterogeneous priors.² Morris (1995) provides a survey of the considerable literature which followed. He argues that in order to avoid using subjective heterogeneous priors to justify any result *ex post* given the sensitivity of many economic models to assumptions over beliefs, it is important to identify systematic regularities in any observed biases to discipline the use of

*We would like to thank Robert Akerlof, Gordon Brown, Dirk Engelmann, Peter Hammond, Graham Loomes, Heracles Polemarchakis, Andrew Oswald, Francesco Squintani and John Stovall for useful comments and Alex Dobson, Stephen Lovelady, Andrew Taylor and Neil Gatti for expert research assistance. Funding was provided by the Leverhulme Trust.

¹While this remark was certainly not actually made by Marie Antoinette, it has been used to epitomize the apparent inability of the ruling classes in pre-revolutionary France to appreciate the difficulties of those significantly poorer than themselves, possibly because of a lack of concern, empathy, or a self-centred set of perceptions.

²For example, Geanakoplos (1989) is a classic reference.

heterogeneous prior beliefs.³ This provides motivation for our study: to provide empirical support for the existence of biased beliefs, catalog systematic regularities and better understand the underlying causes.

To that end we design an incentivized experiment in order to assess the role of information availability in shaping potentially biased beliefs. We considered attitudes (political stance and happiness), choices (mobile phone purchases and hypothetical restaurant choices, which also touches on the role of deference in decision-making) and physical characteristics (height and weight). Characteristics like height and weight *among students in the same university* are perhaps more observable in day-to-day life than say political stance, and so we can check the importance of observability for the degree of bias. Another special feature of our experiment is that (unlike the large part of the economic literature on biased beliefs) we explicitly avoid analysing performance-related beliefs in order to avoid mixing belief-formation with overconfidence which can provide a confounding influence (see Grossman and Owens, 2012, and Burks et al., 2013). Our questions are generally of two types. The first type of questions ask what percentage of peers a subject thinks lie below him or her in a cumulative distribution. In the second type of questions we ask our subjects to estimate the averages for the characteristic, attitude or behavior from among their peers. In all but the hypothetical restaurant choice problem, beliefs are formed through experiences in the real world. As a result we have a measure of external validity for our findings.

From a normative point of view, one hypothesis is that perfectly rational individuals with unlimited attention would not display bias for more easily observable characteristics like the weight and height of other people who they are likely to see on a day-to-day basis in the same environment. However, for elements that are not so readily observable like happiness or political stance the best they can do is to use the observation of the realization of the parameter in a sample that they directly observe. Since they are more likely to have acquaintances that are similar to them, they are more likely to observe values of the parameters that are closer to their own.⁴ Our results support this idea to some extent. For instance, individuals in the political fringes perceive themselves to be more representative, as do those who are very happy or sad. Individuals with less popular mobile phone brands think that there are more individuals using the same brand, but the ones with popular mobile brands correctly estimate its distribution. When asked to decide on a tie-breaking rule in the hypothetical restaurant choice problem, those who choose in a given way tend to see their choice as being more popular than is the case.⁵ Put simply, individuals tend to form beliefs that are “self-centred” in the sense that they see themselves as more “average” than is

³As Morris puts it “We should resort to unmodelled heterogeneities of prior beliefs only when we can imagine an origin for the differences in beliefs and we can perform comparative static exercises, comparing the predictions of heterogeneous prior models with alternative explanations.”

⁴This echoes the literature on assortative matching (Becker, 1973) or homophily (McPherson, Smith-Lovin and Cook, 2001; Golub and Jackson, 2011) through which people may associate with those who are similar to themselves.

⁵For example, an individual who opts for a more deferential tie-breaking rule sees others as similarly deferential.

the case. However, albeit to a smaller extent, this self-centered bias *still holds* for characteristics such as weight and height that we might think are more readily observable: taller and heavier individuals think that there are more tall and heavy individuals in the population, similarly shorter and lighter individuals believe the population to be made up of greater numbers of shorter or lighter individuals (respectively) than in reality.

The fact that a self-centered bias still holds for more observable characteristics perhaps reflects the fact that individuals are *not* endowed with an unlimited amount of attention, a point emphasized elsewhere (e.g. DellaVigna 2009). Hence they may put *undue* weight on easily available data (the so-called *availability heuristic*, Tversky and Kahneman, 1973). Their initial estimate is their own individual value and then they update by using easily attainable data from others to adjust.⁶ Using data deriving from the same experiment we will argue that this is plausible, but it is unlikely to represent the only explanation. If subjects make inferences by using information from their own similar peers, when they are asked to estimate averages of the characteristics, these estimates should be increasing in the subject’s own characteristics. Furthermore, the average individuals, who should draw observations from other average individuals, should not display biased perception. Data from our experiment seem to show this pattern only in part. For example, male subjects tend to perceive average height to be systematically lower than it actually is and subjects in general tend to perceive a lower level of general happiness. We argue in the discussion section below that a possible explanation of this pattern is that subjects’ beliefs are self-centred not only because of limited attention, but possibly because some biases might be “self-serving” since individuals in some instances may not wish to think of themselves as unusual or extreme.⁷ As Carrillo and Mariotti (2000) and Benabou and Tirole (2002) argue, self-serving beliefs may be due to some form of “strategic ignorance” (or “self-serving bias”) and might be “rational” in the sense that it is utility-maximizing to self-delude.⁸ Finally, our data seem to rule out that the effect is due to salience of individuals with extreme characteristics (such as being very tall, or politically at the far right or left). We asked subjects to estimate the averages for subjects belonging to the top and bottom 10% of the different distributions. Any effect due to salience should be characterised by individuals systematically overestimating the averages at the top and underestimating the averages at the bottom

⁶This process is called anchoring in cognitive psychology (Tversky and Kahneman, 1974). The failure to realize that inferences are biased by a non-representative sample can be categorized as irrational and there is a long literature which follows this line of thinking.

⁷There is nothing a priori to rule out individuals wanting to delude themselves into thinking they are unusual or different so this is very much an empirical question, which provides further justification for the current paper.

⁸Of related interest, this literature has shown that biased beliefs may not be inconsistent with a rational process of Bayesian updating, but merely reflective of the imperfect information available to the observers: Benoit and Dubra (2011) show that if individuals have imperfect knowledge of their own ability, even individuals performing correct Bayesian updating starting from a common prior may report what seems to be an overconfident belief.

of the distribution. We do observe this pattern from our data.

To relate our findings back to other work within Economics, we can distinguish three main forms of bias: individuals tend to be *overconfident* about their own abilities (e.g. Moore et al., 2008; Burks et al., 2013); the so-called *law of small numbers*, when individuals expect random draws to be excessively representative of the distribution from which they are drawn (e.g. Clotfelter and Cook, 1993; Rabin, 2002); and *projection bias*, where individuals' expected future utility is too close to current utility (e.g. Loewenstein, O'Donoghue, and Rabin, 2003). Added to this, the literature within Psychology has already emphasized that individuals may suffer from "false consensus" (Ross, Greene and House, 1977) when individuals are asked to indicate their opinion (typically to a yes/no, agree/disagree question). They are then asked to estimate the percentage of their peers who would respond one way or the other and their estimate of consensus for their own position typically exceeds the estimate made by those who endorsed the opposite position.⁹ Recently, economists have contributed to this literature: Engelmann and Strobel (2012) show that in a laboratory experiment false consensus arises only when information is not readily available, but requires some processing by individuals in order to be of use. Burtler, Giuliano and Guiso (2013) show that more trustworthy individuals form more optimistic trust beliefs. However, in our work, unlike in the false consensus literature, we aim to assess the effect of incomplete information in shaping beliefs and to investigate the reason for the resulting biases. Furthermore, there is a fundamental point of departure between our paper and the false consensus literature: while false consensus often relates to matters of opinion, in the current paper we are concerned with matters of fact, and hence we investigate how beliefs differ with respect to the truth and not with respect to the opinion of others.¹⁰ Another concept usually associated with false consensus is "assumed similarity" (Cronbach, 1955). This is generally measured as the absolute difference between the position attributed to oneself and that ascribed to a benchmark individual known to everyone, and generally each individual tends to position the target closer to themselves. This is different from the issues addressed in our work in two important ways. Firstly, normal tests of the assumed similarity bias do not consider the accuracy of perception (i.e. whether they are right to assume similarity), the aim of the current paper, since they are instead interested in comparing individuals' relative perceptions. Second, they do not consider how

⁹See page 612 in the Handbook of Experimental Economics (Kagel and Roth, 1995) for more on false consensus.

¹⁰In one seminal paper on false consensus, Marks and Miller (1987) note: "...The false consensus hypothesis has no direct bearing on whether subjects will overestimate, underestimate, or accurately estimate the actual consensus for their own behavior." For example, consider the view of a right-wing individual who is asked what percentage of people are right-wing. She might say 60%. A left-wing person who is asked the same questions might say 50%. The "false consensus bias" is concerned with the difference between 60% and 50%, while our concern is the difference between 60% and the true percentage of right-wingers and how this difference between beliefs and the truth changes as the type of individual changes. Notice that we can derive one observation from false consensus which closely links to our work: if individuals do differ in their opinions then it is immediately apparent that at least one person must have an opinion that is at odds with the truth. However, this does not help us to see which individual is more likely have such a view.

assumptions of similarity might change across the distribution of the object of interest. In our work a key feature is how perceptions change as characteristics or choices change: for instance whether shorter or taller individuals have more or less accurate beliefs.

The self-centered bias identified in this paper has important implications for economics, for example informational herding theory and auction theory, and for economic modelling more generally. In the latter case there is the ongoing issue of how to model irrational expectations, and certainly thinking about beliefs as self-centred might be a reasonable starting place. There are also direct implications for policies directed towards tackling problems that are partly based on beliefs about where people lie in distributions. A good example would be policies designed to tackle obesity. In a world in which those who are very overweight do not perceive themselves as overweight, stressing the harm of being overweight should perhaps take second place to educating people about their own position in the distribution. Another example might be the behavior of voters who fail to recognize how extreme certain political views might be and can perhaps isolate themselves from more mainstream views. We discuss many applications of both the general and specific type and corresponding policy ramifications in section 5.

2 Experimental Design and Key Variables

A key feature of our design is the merger of field elements in a laboratory-based design by conducting a laboratory experiment which draws on our subjects' real-world experiences. We want to make use of real-world choices and characteristics as far as possible to lend external validity to our findings. We also want to be able to exert as much control over our subjects as possible and to incentivize them for accuracy in their estimations which makes a laboratory ideal.

Our data was collected using a series of computerized tasks and questions presented in a controlled experiment at the University of Warwick.¹¹ The experiment was conducted in a laboratory, however many of the choices and characteristics are drawn from the real-world experiences of our participants as well as information given to them in the laboratory. The participants were 154 students recruited from the university-wide experimental pool of over 1500 subjects.¹²

¹¹The experiment took place in 19 sessions with about 8 students per session, and was conducted on 27 May, 30 May and 29 June 2011. There was also an earlier non-incentivized pilot experiment which consisted of 120 participants drawn from the same experimental pool, held on 17 March, 5 May and 11 May, 2010. The main results for this paper will be drawn from the fully-incentivized experiment, though the data from the pilot study will be used when calculating the average height, weight, happiness and political stance of the Warwick student body.

¹²The University of Warwick keeps a register of those available for use as experimental participants and a research assistant (rather than the experimenters) drew from this large pool of potential applicants on a random basis. Participants were recruited without any knowledge of the nature of the experiment. The times and dates of the sessions were varied to avoid discriminating against participants from any demographic, and in the event we had large variety in terms of subject, year-group and gender.

Subjects were given a £2.50 show-up fee, plus a bonus of £5 pounds if a randomly drawn answer was within 10% of the correct answer in rounds 2 to 6 as described below. For example, if participants were asked to state the average height of the student body in Warwick and this was the randomly allocated bonus question they received a £5 bonus if and only if their answer was within 10% of the true average.¹³ The payment scheme was fully transparent to all participants and highlighted in the instructions during the experiment. No participant was allowed to participate more than once.¹⁴ The experiment itself typically lasted 20 minutes and the average payment was a little over £5, producing an hourly rate of around £15 (equivalent to around 21.50 US dollars or 19 euros). Once each round was completed participants could not go back and change earlier answers, nor did they know the content of later rounds upon entering answers to earlier rounds. This was important as it prevented any attempt to retroactively alter their answers to make winning the bonus payment easier.

The experimental time-line is summarized below but a full transcript of the on-screen instructions is provided in the Supplementary Information.

In the first round participants are asked to report their gender (1a), height (1b), weight (1c), happiness on a 7 point Likert scale (1d), political beliefs on a 7 point Likert scale (1e), and current brand of mobile phone (1f). They are also given a hypothetical restaurant choice as follows: “Imagine that you have to decide between two restaurants in which to have dinner alone. They are called restaurant A and B. You have some private information that A is better, but you know that an equally well-informed colleague has information suggesting that B is better. Would you choose to eat at A, B or are you indifferent?” (1g). This gives participants the opportunity to display a degree of deference to others information or towards their own. (1g) might be of special interest to those interested in rational herding and informational cascade literature (see Banerjee, 1992, and Bikhchandani, Hirshleifer and Welch, 1992) and essentially asks what participants would do in a situation of theoretical indifference when processing information. There is no clear right or wrong answer to (1g) though “A” points to a measure of confidence in the participant’s private signal over that of their colleague, whereas “B” perhaps implies a measure of deference towards others (or a lack of confidence in the individual’s own signal). In the second round participants were then asked to report the percentage of students at Warwick they thought were less happy than they were (2a), less right-wing (2b), shorter (2c) and lighter (2d).¹⁵ For (2c) and (2d) they

¹³In the same vein, Engelman and Strobel (2012) use an incentive scheme to reveal accurate beliefs; they pay a bonus of 5 euro to guess the number of subjects choosing a lottery. For our purposes the true average was based on the numbers generated within this experiment and from the earlier non-incentivized pilot experiment. For one question, denoted (5e) below the scheme was changed slightly as participants had to select from an interval and so they were told that an answer in the correct interval or the one to either side would be sufficient to win the prize.

¹⁴Participation in the pilot experiment also ruled out participation in the full experiment.

¹⁵Note that for this question and in the following questions that specifically reference “Warwick students” we used the information collected during the experimental sessions and during the pilot sessions to form our measure of the true value.

were asked to consider only their own gender. In (2e) they were asked to consider the mobile phone brand listed in round 1 and asked what percentage of students at Warwick (again, not just in their session) they thought also used the same brand of mobile phone as their main mobile phone. For (2f) they were asked to consider the hypothetical restaurant choice (and any implied deference) and report what percentage of their fellow Warwick students they thought chose the same answer that they did (they were reminded of the entirety of the question and the possible answers). In round 3 they were asked to report the average height for someone in the 10% tallest Warwick students of their gender (3a), the average weight for someone who is in the 10% heaviest Warwick students of their gender (3b), the average happiness for someone who is in the 10% happiest students at Warwick (3c), and the average political belief for someone who is in the 10% most right-wing students at Warwick (3d). For (3c) and (3d) they were asked to use a 7-point Likert scale as before. Round 4 was phrased identically to round 3 except that in each case in the four questions they were asked to report the average for the 10% shortest (4a), 10% lightest (4b), 10% most sad (4c) and 10% most left-wing (4d), again for the population of students at Warwick, using a 7-point Likert scale for (4c) and (4d), and considering only their own gender for (4a) and (4b). Round 5 focused on overall averages rather than extremes in the distribution. They were asked to report the average height (5a) and weight (5b) for a Warwick student of their gender, and using a 7-point Likert scale the average happiness (5c) and political belief (5d) for a Warwick student. In question (5e) they were asked to estimate the percentage of their fellow Warwick students who used each of a selection of mobile phone brands. They were presented with a tabulated list of the most popular brands in the UK, and they were informed that the list was presented in alphabetical order (except for the “other” category which was presented last). They were asked to include an entry for every brand (including “other”). For (5f) they were asked again about the hypothetical restaurant choice: “Think again about the restaurant question you were asked earlier in the session. To remind you, you had to decide between two restaurants in which to have dinner alone. They were called restaurant A and B. You had some private information that A is better, but you knew that an equally well-informed colleague had information suggesting that B was better. What percentage of your fellow Warwick students do you think would have chosen to eat at restaurant A if they were asked the same question? Remember that the other options were indifferent and B.” For round 6 the participants were asked to answer a single question designed as a check on their ability to understand and manipulate expectations and probability: “Consider the following gamble. You have a 20% chance of winning £100, a 40% chance of winning £10 and a 40% chance of winning £0. If you played this gamble many times what would you expect to be your average winnings per gamble? (in pounds)” Round 7 is a final questionnaire and, as is conventional,

was not incentivized (participants were informed that the incentivized part of the experiment had ended) since there was no way of checking right or wrong answers. They were asked to report their age (7a), nationality (7b), degree subject (7c), whether they studied mathematics up to their final year at school (7d) and also comment on their methods, if any, during the incentivized parts of the experiment (7e).

3 Results

Before moving on to the analysis we present the main variables in table 1 below. The variable “Happiness” is coded from completely sad (1) to completely happy (7) and taken from the answers to question (1d), and “Political Stance”, from extreme left (1) to extreme right (7), taken from question (1e). “Weight” is converted to kg from the answers in question (1c) and “Height” to cm from the answers in question (1b). The table allows us to see not just the range within our sample but also should give an idea of the variance across difference characteristics. The numbers also indicate that our sample is not unusual relative to the UK or English average.¹⁶ There are 154 subjects participating to the experiment, 84 males and 70 females. We prefer not to force the answers to avoid random or false answers from subjects who did not want to truthfully report some measure. There are 3 subjects that did not report their weight (1 male and 2 female), 7 subjects who did not report their height (4 male and 3 female). These subjects guessed their place in the distribution and have been regularly paid at the end. This small attrition is not unusual on the experiment involving answers about personal questions and given the small number for each category, it is highly unlikely that these missing answers have sensibly biased our results that, as it will be clear below, hold with a high level of confidence.

In this section we first present an analysis of individuals’ beliefs about their own positions in the true distribution, starting first with a model. Thereafter we analyze estimated averages. In both cases our findings indicate that beliefs about position in the distribution and about averages are a function of individual’s own position in the distribution.

3.1 Modeling Beliefs

Let Θ be the set of characteristics, attitudes and choices considered in this paper, for example height, reported happiness or the market share of a brand of mobile phone. $\theta \in \Theta$ is a particular characteristic, attitude or choice, for example, height. We index the value of each characteristic for each individual i , so

¹⁶For example, Moody (2013) indicates that for England, average male height is 175.3cm for adults over 16, and 177.8cm for young adults between 25 and 34, while for females it is 161.9cm for over 16s and 164.5cm for females between 25 and 34.

θ_i might be individual i 's height. Let $F(\theta)$ be the cumulative distribution of characteristic θ . In a mild abuse of notation, denote $F(\theta_i)$ as individual i 's true position in the cumulative distribution and then $E_i(F(\theta_i))$ is i 's belief about his position in the cumulative distribution.¹⁷ Given the nature of incentives in the experimental design a typical subject i solves the following programme:

$$\text{Min}_{E_i(F(\theta_i))} \{E_i(F(\theta_i)) - F(\theta_i)\} \quad (1)$$

Technically the programme to be solved is slightly more complex because we reward any guess that is within 10% of the truth. This implies that it is reasonable for a subject to guess just below their true beliefs, especially if they exhibit risk aversion. However, this bias would generate a best guess that is very close to the truth and some distance from the reported values that we find and therefore cannot explain the results in this paper. Truth-revelation mechanisms do exist but the significant increase in complexity (and the resultant concern that subjects might not fully understand the payment scheme) would far outweigh any benefit.

We assume that individuals form their expectations by drawing from a sample M_i from the total population N . However we cannot rule out the possibility that the sample of size M_i may be biased. In particular it may be that the sample is taken from members of the population who are similar to the individual i . An unbiased estimator would be

$$E_i(F(\theta_i)) = F(\theta_i) + \text{error}_i \quad (2)$$

where error_i is i.i.d., with mean 0 across the population. In order to assess the existence of a systematic bias, we will therefore consider the model:

$$E_i(F(\theta_i)) = G(F(\theta_i)) + \sigma(\theta_i)\epsilon_i, \quad (3)$$

Where $\sigma(\theta_i)$ is a general function to allow for heteroscedasticity with respect to θ_i , and ϵ_i is a white noise error.¹⁸

From the data, we can observe the beliefs $E_i(F(\theta_i))$, the real distribution $F(\theta_i)$ hence we estimate $G(F(\theta_i))$ both non parametrically using local polynomial smoothing¹⁹ and by making parametric

¹⁷We preferred to ask “less than”, rather than “at least as” because we judged it a more natural question that individuals are likely to have faced in their everyday life. We calculated the CDFs using the same definition.

¹⁸Note that we are conceptualizing what is likely to be a subconscious belief-formation process in a form that is general enough to encompass several suggested rationales for our findings given in section 4 so for instance the error term can include both belief-formation errors and reporting errors.

¹⁹Kernel-Weighted Local polynomial smoothing involves fitting the response to a polynomial form of the regressor via

assumptions on this functional form.

This formulation allows us to examine directly the difference between $G(F(\theta_i))$ and the true distribution, $F(\theta_i)$ which represents the bias made by each individual, allowing for non-linearities in beliefs as we change the own values of each individual. One way to add structure to this bias might be to define $G(F(\theta_n))$ as $G_{M_i}(F(\theta_n))$ with the G function acting as a way to choose a subsample $M_i \subseteq N$ where the statistic $F(\theta_n)$ is defined. In particular we might then consider $M_i = N$, when information is perfect and individuals have unlimited attention, hence $G_N(F(\theta_i)) = F(\theta_i)$, with no implied bias. Alternatively $M_i = N_i$, with N_i = individuals belonging to a group of similar peers when information is not perfect and outsiders are not observable and/or when individuals have limited attention and can only remember other individuals who can observe more often (thus if there is homophily, individuals in N_i are more similar to i). Or we might consider $M_i = M_i^*$, with M_i^* as a sample selectively chosen by the individual, when memory is self-serving in the sense of Benabou-Tirole (2002).²⁰ We purposefully leave the potential nature of the bias open to allow us to look for (and hopefully differentiate between) several different sources, which we will return to in section 4. In the following subsection we present the empirical evidence with respect to the characteristics weight, height, happiness and political stance. We will also analyze the beliefs about mobile distribution and deference in information processing. These last are slightly different because $F(\theta_i)$ represent frequencies and not CDFs.

3.2 Empirical Evidence on Perceived Distributions

Given the nature of the data, we present the results in three separate groups. The CDFs of the observable characteristics, height and weight, are in figure 1; the CDFs of the unobservable characteristics, happiness and political stance, are in figure 2.²¹ Finally, the frequencies of answers to the restaurant choice question (1g) and the mobile choice question (1f) can be read from the histogram in figure 3.

We start by analyzing weight, height, happiness and political stance. From the right panels of figures 1 and 2. We can observe the real CDF, $F(\theta_i)$ and the scatter-plot of the perceived positions, $G(F(\theta_i))$ with their locally weighted scatter-plot smoothing or “Lowess”.²² The scatters in the right panels of

locally weighted least squares. In the Kernel-weighted regression, $G(F(\theta_n))$ is calculated without assuming a functional form, as a constant term of a regression weighted by the kernel function of $E_i(F(\theta_i))$ on the polynomial terms $G(F(\theta_n)) - G(F(\theta_i))$, $(G(F(\theta_n)) - G(F(\theta_i)))^2, \dots, (G(F(\theta_n)) - G(F(\theta_i)))^p$, for each point $G(F(\theta_n))$. The definitive reference is Fan and Gijbels (1996), see also the Stata 12 base reference manual, p. 1001.

²⁰ Along similar lines, see Compte and Postlewaite (2004) who motivate their paper with the story of a lawyer who observe a full history of his successes and failures but who may dismiss his losses as stemming from biases in the judicial system and instead act as though he has only ever won each case. They then go on to show how equilibria in which this sort of reasoning takes place can exist in a model of rational overconfidence.

²¹ When calculating the population averages we also used the 120 data from the pilot experiment to increase the size of the sample to 274.

²² This displays for each value of the independent variable, θ_i a smoothed value of the dependent variable, $G(F(\theta_i))$. The

figure 1 represent the real and the perceived CDFs in θ_i for each subject i , θ_i for height and weight, and in figure 2 the same observations for political stance and happiness. The dashed line represents the estimated function $G(F(\theta_i))$ with the 95% confidence interval; In order to compare $G(F(\theta_i))$ and $F(\theta_i)$, we also plot the 45 percent degree line. There is a clear pattern: individuals at the extreme tend to overestimate the number of individuals who are equal to or more extreme than themselves, while those at the center of the distribution seem better informed. This is true for happiness and political stance, female and male weight, and male height. For political stance and happiness we note that the two extreme numbers (1 and 7) seem to revert to the 45 degree line. This makes sense since these two characteristics are bounded above and below, therefore extreme individuals will be able to calculate more accurately their position (and have lower scope for self-deception). This effect appears even more clearly in the parametric analysis to follow.

We analyse parametrically how $G(F(\theta_i))$ is influenced by $F(\theta_i)$ for happiness, political stance, height and weight. We use a linear econometric model because as we will see it will imply a simple interpretation of the coefficients. The shape of the polynomial interpolations in the left panels of figure 1 and 2 suggest that this is a good approximation. We always include a male dummy, in order to control for heterogeneity between genders. We therefore estimate the following model:

$$E_i = \alpha + \beta F(\theta_i) + \delta male_i + \epsilon_i \quad (4)$$

where E_i is the individual estimation of $F(\theta_i)$ and ϵ_i is an individual specific error term. When we consider height and weight distributions we introduce the interacted term $\gamma male_i F(\theta_i)$ in the model 4. The male and female variables have very different influence on both these characteristics, hence we allow for the possibility that gender has different effects upon their respective perceptions.

As we can observe from the right panels of figures 1 and 2, there is some heteroskedasticity in E_i given that individuals at the extremes can be more precise in their estimations. For this reason we use robust standard errors. The linear model we assume, implies that a positive and significant constant term, $\alpha > 0$, would signal an overestimation of sadder, more left-wing, shorter and lighter individuals and a coefficient $\beta < 1$ would signal a overestimation of happier, more right-wing, taller and heavier. A β closer to one would signal a more accurate estimation on average.²³

smoothed values are obtained by running a linear regression using only the data $(x_i; y_i)$ and a subset of the data with the x values close to x_i . Data are weighted so that the central point $(x_i; y_i)$ receives the highest weight and points that are farther away receive smaller weight. The estimated regression line is then used to predict the smoothed value for y_i only; a separate weighted regression is performed for every point in the data.

²³A possibility is that individuals that know they are likely to be at the top/bottom of the unbounded distributions of weight and height might strategically declare to be in a lower quantile because we allow a 10% error band. As a further

Considering table 2, we find all coefficients β less than 1 ($p - value < 0.01$), the closest to unity is the coefficient related to height, then in decreasing order for weight, happiness and political stance. The β coefficients for male and female are remarkably similar for weight and height.²⁴ We also find positive and significant intercept, α , with the only exception being the coefficient on female height, which is not significantly different from 0. This indicates that females towards the lowest extremes are the most accurate in their perceptions about this characteristic.

Mobile phone ownership and attitudes in the hypothetical restaurant choice are not ordered variables, it is therefore impossible to determine a perceived CDF and a real CDF. We therefore collected data on beliefs about frequencies and compared these with the true frequencies. Accordingly, we define our $F(\theta_i)$, as the true frequency of the choice θ_i and $G(F(\theta_i))$, the beliefs' function. The frequency of the different brands and the hypothetical restaurant choices can be observed in right panel of figure 3. While all 154 subjects reported the answer to the restaurant choices, only 141 subjects declared to have a mobile (hence mentioned its brand).

In the left panels of figure 3 we can compare the histogram representing the distribution on mobile phones and attitudes in the hypothetical restaurant choice question, $F(\theta_i)$, with the subjective beliefs E_i , where the different θ_i have been ordered increasing with their frequencies. Although we are comparing now frequencies rather than CDFs, we note a similar pattern to the one emphasized in figures 1 and 2. The Lowess tends to stay above the real frequencies for less popular brands and for less frequent tendencies in the restaurant choice question.

As before we estimate the model 4, and present the estimated $G(F(\theta_i))$, in the right panels of figure 3. The reading of the panels is slightly different from before, with the average subjective beliefs above the real level for less frequent θ and are non-significantly different for more frequent θ . The interpretation of the results are the same as with happiness, political stance, height and weight: individuals with more common brands or with the most common attitude in the restaurant question have beliefs that are on average correct, on the contrary individuals making less common choices overestimate the frequency of their choices.

In table 3, we estimate parametrically $G(F(\theta_i))$ for the mobile phone distribution question (2e) and attitudes towards information processing from the restaurant question (2f). As before, the right panels of figure figure 3 seem to suggest the presence of some heteroskedasticity, and we again use robust standard

check that this is not biasing our results, we also run the same regressions that are presented in table 2 by excluding the top and bottom 10% of the distribution, obtaining almost identical results. These regressions are reported in table A.1 in the appendix.

²⁴As a robustness check, we also run the happiness and political stance regressions allowing for two separate coefficients for male and female, the results are qualitatively similar, the female have a slightly higher β coefficient on happiness and a lower β coefficient for political stance. These regressions are reported in table 4.

errors. From table 3 we note that the β coefficient is less than unity and the intercept α is more than 0; hence individuals with a less popular brand tend to significantly overstate the frequency of other individuals using the same brand. This confirms the main finding emphasized by figure 3.²⁵

In essence subjects towards the extremes of the distributions think there are more individuals in the same position (or indeed in a more extreme position) than themselves, while more average subjects tend to have more correct beliefs. So, a very tall person really does perceive him or herself as more “normal”. Perhaps more economically significant, someone who has purchased a less popular mobile phone believes it to be more popular than is the case. Note that this bias is true for different frames. Either when subject are asked about their position and when their asked about the frequency of their choices.

One remarkable point that stands out is the ubiquitous nature of this bias: only for females who are among the shortest and only in the parametric analysis is there any deviation from the simple rule that those at the extremes do not see themselves as being as extreme as they truly are.

More formally, we find that individuals with a θ_i closer to the average tend to estimate better their position than those with a θ_i closer to the extremes of the distribution. Denote E_i as the expectation operator for individual i . We can then express this as a simple conjecture: for each $\theta \in \Theta$, $\exists \theta^*$ s.t. $\theta_i < \theta^* \Rightarrow E_i(F(\theta_i)) > F(\theta_i)$ and $\theta_i > \theta^* \Rightarrow E_i(F(\theta_i)) < F(\theta_i)$. Put simply, we can consider a point in the cumulative distribution θ^* such that individual’s with value $\theta_i < \theta^*$ believe they are positioned higher in the distribution than is the case, and for $\theta_i > \theta^*$ they believe they are lower.

This is a form of mean-reversion of beliefs. Since beliefs seem too heavily dependent on θ_i we label them “self-centered”.

4 Discussion: What Generates Self-centered Beliefs?

The above analysis suggests that we can rule out the possibility that individuals are fully Bayesian and endowed with unlimited attention, something that we can define as a perfect “statistical model” of learning. When incentivized subjects are asked to estimate their position in the distribution of height and weight it seems likely that they would refer to causal observation gathered in their everyday lives. Consider for instance, height, and how many individuals our typical subject might have seen in the few hours prior to the experimental session alone. In a large campus university with several thousand students, an individual subject is likely to have passed several hundred other students just getting from

²⁵We also note a slight tendency for the owners of the very smallest market share brands to better understand their own-brand market-share as compared with those with slightly higher market share brands. This could be due to the fact that the owners of these least popular brands were looking for niche products with specific features and so spent more time and effort on research or may see themselves as different.

her accommodation to the laboratory. If an individual, when asked to assess his or her position in the dimension of an easily observable characteristic, considers only his or her group of close acquaintances, it seems much more likely that they might be (rationally or irrationally) anchoring themselves to this group rather than having only seen individuals from this group.

There are many possible routes through which the self-centered beliefs might emerge. Here we discuss how well they fit our data. We can broadly distinguish three different theories: (i) limited attention; (ii) self-serving beliefs; and (iii) salience of the extremes.

(i) Limited Attention

Consider first an explanation based on limited attention. We might argue that individuals can be perfectly Bayesian in the way they update their beliefs but only use the observation of the realization of the parameter in a limited sample that they observe more often because they have limited attention capacity. The fact that the extremes in the distribution suffers the worst from the bias lends some support to the *availability heuristic* explanation identified by Tversky and Kahneman (1973), according to which individuals put *undue* weight on easily available data, and may then draw biased inferences.²⁶ Closely related is the concept of anchoring in cognitive psychology: individuals base their estimate by first looking at their own individual value and then use the evidence from others to adjust (Tversky and Kahneman, 1974). The failure to realize that inferences are biased by a non-representative sample because of anchoring can be categorized as irrational.

We can test the plausibility of this explanation in figure 4, where we plotted the beliefs about the averages perceived by each subject against the individual characteristic of the same subject, and a line representing the real average, calculated by using the subjects in our sample. If when asked about the averages, individuals only sample their close peers and if individuals associate with similar peers, we should observe a positive association and that the regression line should cross the real average line close to the middle, so that average subjects that match with other average subjects make unbiased estimation.

Figure 4 suggests that a limited attention explanation can explain self-centred beliefs to some extent, but it is unlikely to be the only determinant. We can observe that for weight, both male and female subjects in the middle are the most correct, consistently with the limited attention story. Furthermore, we can observe a positive relationship between the characteristics and the perception in all but in female height. For height, male subjects in the upper end of the scale seem the most precise, which is consistent with the results in figure 1, where we can observe they are also the ones that more correctly estimate

²⁶Linked to this, is the idea that observing others is one of the cheapest ways to acquire information and is the key form of learning in the social learning and herding literatures (see Banerjee, 1992, and Bikhchandani, Hirshleifer and Welch, 1992).

their positions in the distribution.

Figure 5 examines our subjects reported beliefs and true values of the average, as well as the top and bottom deciles for happiness, political stance, male/female weight and male/female height. Looking at the bottom right-hand panel of figure 5, consistently with the above observation, we also note that in general male subjects significantly underestimate real average height. For female height the relationship is negative and insignificant, which poses an interesting puzzle, and female subjects seem also to slightly overestimate the average. Furthermore, average happy subjects seem to slightly underestimate the level of happiness, which seems to be true also in general, as we can note from the top left-hand panel of figure 5, where average happiness is significantly below the value reported by the subjects. Finally, we note that in political stance subjects believe the true distribution to be slightly more to the right than is the case in our data. In what follows next, we will try to explain these anomalies.

(ii) Self-Serving Beliefs

When forming beliefs, focussing on peers that are similar can be a strategy rather than due to a limited amount of attention. Costly learning can explain this form of deception since we can argue that where falsely believing yourself to be in a particular position is useful, it might cost more (in terms of final utility) to learn otherwise.²⁷ To explain why our subjects fail to learn the truth, we might even appeal to the optimal experimentation literature, beginning with Robbins (1952) and brought to economics by Rothschild (1974). This literature demonstrates that optimizing decisions while simultaneously attempting to learn can result in learning stopping before the truth is known. Since our subjects need to make decisions on a daily basis using the beliefs they form, this process of decision-making and learning in a simultaneous setting may be appropriate. In relation to the concept of “self-deception”, Carrillo and Mariotti (2000) and Benabou and Tirole (2002) provide numerous mechanisms (including self-delusion and memory manipulation) to foster rational ignorance of the truth (indeed Carrillo and Mariotti call this type of behavior “strategic ignorance”). Grossman (2015) and Grossman and van der Weele (2015) discuss the related concept of “willful ignorance” through which individuals obfuscate signals in order to maintain positive self-image and show that ignorance equilibria exist in which information avoidance is possible.

Therefore, we can argue that individuals, in some cases, may want to report a bias estimation to feel better. This would also explain why average individuals are not always correct in estimating average happiness and height, as we argued above using figure 4. The reason the average subjects tend to underestimate the level of general happiness and male height and overestimate female height might be a

²⁷Coate and Loury (1993) and Farmer and Terrell (1996) look at costly learning in the context of discrimination.

self-serving belief. Short males might want to think that there are more short individuals than in reality, this would also explain why tall males do not typically suffer from this bias (see figure 4, but also figure 1, where we note that the self-centered bias is very small for male subjects in the upper end of the height distribution). At the same time, tall females might want to think that there are more tall females than in reality, but being a small female is not a particular concern (we note from figure 1 that the self-centered bias is very small or non-existing in female subjects in the lower end of the height distribution).²⁸ Finally, we can argue that individuals might want to generally think that their levels of happiness is higher than the average, this can explain the fact that average happy subjects perceive a slightly smaller average level of happiness and that perceived happiness is in general smaller than in reality (as we noted from the top left-hand panel of figure 5).²⁹

(iii) Salience of the Extremes

Notice also that extreme characteristics (for instance, the very tall or extremely left wing) are generally thought to be more salient than those with average values, which could provide another possible bias. We can reasonably rule out this possibility for mobile brands, in this case there is no reason why the extreme, i.e. the owners of a fringe brand, should be more salient than the others, and there is no evidence to support salience in this case. However, for political stance, happiness, weight and height this is a real possibility. To test this we collected data on perceived averages for the potentially most salient (the top and bottom 10%) in rounds 3 and 4 of the (incentivized) experiment as described in section 2. For example, we asked subjects in question (2b) to report their average for the top 10% heaviest individuals (of their gender). In figure 5 we plot bar graphs of the perceived averages in the top 10%, in the bottom 10% of the distributions and the averages in the entire distributions together with the values estimates using our sample, for Political Stance, Happiness, Weight and Height. If there is an effect due to the salience of the extremes we should observe that the perceived averages in the top should be constantly higher than the “real” values and the ones in the bottom should be constantly lower than the real values.

From figure 5, we cannot observe this pattern and these estimations seem in general remarkably accurate. The only cases consistent with the effect due to salience is: the perceived average in the top 10%, which is higher than the real values in female weight; and in the bottom 10% of female height, we observe that the perceived value is smaller than the real value. However, for male height this pattern is reversed: the real average of the top 10% is actually higher than the perceived one, and for the other 9

²⁸Consistently with this results, Oswald (2008) finds that perceived height is increasing but concave in real height.

²⁹The argument that belief in above average happiness is self-serving can be supported by the “illusory superiority” bias, well-known in social psychology which has been found to explain a general overestimate of individuals’ own-qualities in everything from IQ to happiness in social situations. See for instance Buunk (2001).

cases there is no significant difference between the real and the perceived values. From figure 5, we can then argue against the salience explanation for self-centred beliefs.

Closely related to salience is the evidence that subjective probabilities are inverse S-shaped with respect to true probabilities, or more simply that individuals overweight small probabilities and underweight large probabilities.³⁰ Interestingly enough, our findings might then provide evidence that individuals perceive distributions like probabilities, therefore we could re-interpret the questions we ask in round 2 slightly differently. For example consider being asked what percentage of people might be shorter than yourself. This could be re-interpreted as asking for the probability that a random individual is shorter, and so would be subject to the inverse S-shape relationship. We would then expect to see lower probability events (for instance the chance that someone who is very short being drawn) being overestimated, while seeing higher probability events (for instance the chance that someone is taller than a short person being drawn) being underestimated. This would imply a cumulative distribution function of beliefs biased in a similar way to those we derive from our data: those at extremes would tend to believe that the probability that some individuals are even more extreme than they are, is greater than is the case, pushing themselves down the distribution towards the mean. We cannot entirely rule out this explanation, however this would not explain why perceived averages are correlated in 5 cases of 6 to individuals' own values, as we can observe from figure 4 and so would be a partial explanation at best.

5 Applications

The implications and applications that stem from our work are of three types. Firstly the self-centered bias identified in this paper has important implications for economic modeling generally and the interpretation of data derived from economic models. Secondly, there are implications for certain theories, particularly those which have beliefs at their core, such as informational herding or auction theory. Finally, the results from the specific cases we examine, such as height and weight, can also be used directly, for instance when designing government policy to tackle obesity.

To start with the first type of implication, Manski (2004) highlights how departures from rational expectations can leave economists in some difficulty when seeking to identify the correct model through revealed preference. The bias identified in this paper represents such a departure and the ubiquitous nature of our findings suggests that in many behavioral or choice contexts a self-centered bias may be playing a role in forming expectations. It is therefore difficult to know whether any decision is taken

³⁰There is a large literature on this, but a summary can be found in Camerer, Loewenstein and Rabin (2003).

through an erroneous belief or through the reasonable use of private information or preferences. On that basis it is hard to draw conclusions concerning which model is correct from choices when the choices themselves may have come from more than one competing model and the paradigm of revealed preference becomes suspect. To be more specific, consider the example of the ultimatum game used by Manski. In the game the sharing rule is determined by the proposer's expectations concerning the respondent's behavior. Manski shows that if one departs from the assumption that the objective probabilities of the respondent's behavior are known, data cannot identify a single model for play in the game. The self-centered bias identified in this paper plays a similar confounding role since it is plausible that the proposer forms his expectation on the basis of his perception of the distribution of preference for fairness which may be self-centered.

We can also consider lessons that come from our findings for economics and policy-makers. We would warn policy-makers and survey-designers that the assumption that beliefs are on average correct, even concerning such seemingly straightforward characteristics as height or weight, seems woefully inadequate in any context where a self-centered bias may emerge and our findings indicate that such a context may be far more wide-ranging than has hitherto been considered. For example, bias in the perception of the happiness of others might bias individuals' attitudes towards altruism and redistribution and in turn might bias responses in surveys and should lead us to be more careful in the use of such surveys to guide policy. To restate, for important policy decisions or even in the development of new economic theory it makes sense to think about whether biased beliefs will render a model inaccurate or a policy counter-productive, and if so, it makes further sense to think about how to measure the beliefs of the target population. The solution suggested by Manski (2004) is to make greater use of subjective probabilities in survey-based work and our findings lend empirical support to that recommendation.

Next we can examine the ramifications for models which rely on a clear understanding of how beliefs are formed. Biased perceptions of others' choices can lead to bias in your own choice when that choice is itself a function of the choice of others. This is clear for network goods but extends to any goods or services where quality might be uncertain as discussed in the informational herding literature initiated by Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992). People overstate their own signals or information (believing it is over-representative) so overweight their information when updating. This could be applied to any decision problem but has particular resonance for the experimental herding literature since it can help to explain the well-known finding that individuals overweight their own private signals. This has been a standard feature of sequential herding experiments since Anderson and Holt (1994) and remains a feature of herding experiments in endogenous time (see Sgroi, 2003), herding

with efficient prices (see Cipriani and Guarino, 2005, and Drehmann, Oeschler and Roider, 2005) and herding in endogenous time with efficient prices (see Park and SgROI, 2012). To give another example which links our findings to the social learning literature, in a model of social learning, the perception that individuals are more likely to use their own information than is the case should lead individuals to apply an upward weight to the informational-content of their actions. This would make those in a sequence of decision-makers more likely to overweight the actions of those at the very start of the chain. This fits the story in Guarino and Jehiel (2012) who analyze behavior in a herding model where agents are unable to understand other agents' strategies in their finest detail, and is itself an application of Jehiel's (2005) Analogy-Based Expectation Equilibrium concept.

Economics typically has, at its core, the belief that economic agents are fully aware of the distribution that corresponds to an unbiased distribution surrounding the truth. This may even be a requirement for certain core theories to hold. Take for example the centre-piece of auction theory, the revenue equivalence theorem, which requires that bidders know the true distributions of valuations. The bias identified in this paper would suggest otherwise: bidders are likely to think that other's valuations are more highly correlated with their own than is justified by the truth. The ramifications of such a bias for any policy decision or economic model is considerable: with biases across such a wide range of characteristics and choices, it is easy to see how policies could go badly wrong or models become misspecified.

Finally, we can also draw from our findings to give direct policy ramifications in specific cases. For example, an overweight individual might perceive themselves to be close to average weight and this might leave them content with their diet and exercise regime. That same individual might choose to exercise more or eat more healthily if they realized their true position in the distribution. The same argument applies symmetrically for individuals at the other extreme of the scale. A direct policy ramification would be that one relatively low cost direct government policy intervention to tackle growing obesity is simply to publicize accurate data on average weight, height and BMI together with "ideal" figures.

Bias in the perception of the political beliefs of others could change political actions, including for instance the decision to vote strategically. If a voter thinks that an extreme party close to her views is more representative than is the case, that party may pick up more votes, rather than losing votes because of a perception that they have little chance to secure victory. Ironically bias in perception in this case would actually raise the party's votes and so be partially self-confirming. More generally we might model voting behavior as a function of the proximity of the party in question to a voter's own beliefs and a function of the voter's belief that the party will win (in the vein of a Keynesian beauty contest) and our findings certainly suggest that potential voters, even for quite extreme parties, have unreasonably

high expectations that their favored party is likely to win. Something similar could also be applied to behavior during auctions. If bidders think there is more correlation across valuations than there actually is, this might add a common-value component to an auction which might otherwise be categorized as being entirely private values, or bolster the common value component of an interdependent value auction. One possible outcome might be a scaling down of bids in response to a heightened (and erroneous) belief in the prospect of the winner’s curse coming into effect.³¹ It is possible to imagine a scenario in which several bidders with private values assume a common values component and each scales back their bid resulting in a reduced winning bid and lower revenues for the seller.

The perception that the brand of the good I am using is more common than it is in reality may have important consequences in terms of future consumption choices especially for network goods and status goods, and more generally if the market share is believed to be a good indicator of quality. This is good news for lower market share (niche) items since they are likely to benefit from incorrect beliefs that they are more popular than is the case.

6 Conclusions

Our results highlight a subtle form of self-deception when individuals consider their own characteristics, attitudes and choices: those at the extremes seem to perceive themselves as much closer to average than is the case.³² Our second consistent finding is that estimates concerning averages tend to be a function of individuals’ own values. The biases we uncover are both powerful and ubiquitous, they apply across a variety of different characteristics, choices and parameters, and apply whether they are relatively easy or difficult to observe. Moreover, our experimental design allows us to use the observability of certain characteristics to differentiate between statistical and behavioral theories. Our subject pool are in constant daily contact with their peers and so should have a reasonable recall of the distribution of height and weight and where they stand within that distribution. We cannot rule out that subjects suffer from some form of measurement error when observing others though our findings suggest that this error would have to be biased in a very particular direction and the biases we uncover do not disappear for characteristics or choices with greater or lesser observability. In this way our findings offer strong support to behavioral theories which work irrespective of observability, and would be for instance well in

³¹The “Winners Curse” applies where the winner of an auction factors in the value of the item *conditional on being the winner* which in a common values setting suggests an over-bid absent any correction. A rational bidder will compensate by scaling back a bid.

³²This even includes quite abstract attitudes such as the extent to which others might be deferential when considering the information provided by peers, examined in the case of a hypothetical restaurant choice.

alignment with recent work on self-serving bias.³³ We also present numerous applications which should highlight the importance of our findings which range from the specific (policy applications linked to the characteristics we examine) through to the general (the implications for theoretical modeling in specific areas of economics and more generally).

References

- [1] Baccara, M., and Yariv, L. (2012). “Homophily in peer groups”, forthcoming in *American Economic Journal: Microeconomics*.
- [2] Banerjee, A. V. (1992). “A simple model of herd behavior”, *Quarterly Journal of Economics* 107: 797-817.
- [3] Becker, G. S. (1973). “A theory of marriage: part I”, *Journal of Political Economy* 81: 813-46.
- [4] Benabou, R., and Tirole, J. (2002). “Self-confidence and personal motivation.” *Quarterly Journal of Economics* 117: 871-915.
- [5] Benoît, J-P, and Dubra, J. (2011). “Apparent overconfidence.” *Econometrica* 79: 1591-1625.
- [6] Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992). “A theory of fads, fashion, custom and cultural change as information cascades”, *Journal of Political Economy* 100: 992-1026.
- [7] Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992). “A Theory of Fads, Fashion, Custom and Cultural Change as Information Cascades”, *Journal of Political Economy* 100: 992-1026.
- [8] Burks, S., Carpenter, J., Goette, L., and Rustichini, A. (2013). “Overconfidence is a social signaling bias.” Forthcoming in *The Review of Economic Studies*.
- [9] Butler, J., Giuliano, P. and Guiso L. (2013), Trust, Values and False Consensus. Forthcoming *International Economic Review*.
- [10] Buunk, B.P. (2001). “Perceived superiority of one’s own relationship and perceived prevalence of happy and unhappy relationships.”, *British Journal of Social Psychology* 40(4): 565-574.
- [11] Camerer, C.F., Loewenstein, G., and Rabin, M. (2003). *Advances in Behavioral Economics*. Princeton, New Jersey: Princeton University Press.

³³Self-serving bias might also offer an explanation for the few exceptions to the general rule that people seem to want to appear more average as discussed in section 4.

- [12] Carrillo, J.D., and Mariotti, T. (2000). "Strategic Ignorance as a Self-Disciplining Device", *Review of Economic Studies* 67: 529-544.
- [13] Cipriani, M. and Guarino, A. (2005). "Herd behavior in a laboratory financial market" *American Economic Review* 95, 1427-1443.
- [14] Clotfelter C.T. and Cook, P.J. (1993) "The gambler's fallacy in lottery play" *Management Science* 1993 39:12, 1521-1525.
- [15] Coate, S., and Loury, G.C. (1993). "Will affirmative-action policies eliminate negative stereotypes?", *The American Economic Review* 83: 1220-1240.
- [16] Compte, O., and Postlewaite, A. (2004) "Confidence-enhanced performance", *The American Economic Review* 94: 1536-1557.
- [17] Cronbach, L.J. (1955). "Processes affecting scores on "understanding of others" and "assumed similarity." *Psychological Bulletin* 52: 177-193.
- [18] Currarino, S., Pin, P., and Jackson, M. (2009) "An economic model of friendship: homophily, minorities and segregation", *Econometrica* 77: 1003-10045.
- [19] Drehmann, M., Oeschler J. and Roider, A. (2005). "Herdling and contrarian behavior in financial markets: An internet experiment" *American Economic Review* 95, 1403-1426.
- [20] DellaVigna, S. (2009). "Psychology and Economics: Evidence from the Field", *Journal of Economic Literature*, 47: 315-372.
- [21] Engelmann, D., and Strobel, M. (2012). "Deconstruction and Reconstruction of an Anomaly", *Games and Economic Behavior*, 76(2): 678-689.
- [22] Fan, J., and Gijbels, I. (1996). *Local Polynomial Modelling and Its Applications*. London: Chapman & Hall.
- [23] Farmer, A. and Terrell, D. (1996). "Discrimination, Bayesian updating of employer beliefs and human capital accumulation", *Economic Inquiry* 34: 204-219.
- [24] Geanakoplos, J. (1989). "Game theory without partitions, and applications to speculation and consensus", Cowles Foundation Discussion Paper No.914. Available from <http://cowles.econ.yale.edu>

- [25] Golub, B., and Jackson, M. (2011). “Network structure and the speed of learning: measuring homophily based on its consequences”, *Annals of Economics and Statistics*, Forthcoming.
- [26] Grossman, Z. (2015). “Self-signaling and social-signaling in giving”, *Journal of Economic Behavior & Organization*, 117(1): 26-39.
- [27] Grossman, Z. and Owens, D. (2012). “An unlucky feeling: Overconfidence and noisy feedback”, *Journal of Economic Behavior & Organization*, 84(2): 510-524.
- [28] Grossman, Z. and van der Weele, J. (2015). “Self-image and willful ignorance in social decisions”. Forthcoming in the *Journal of the European Economic Association*.
- [29] Guarino, A., and Jehiel, P. (2012). “Social learning with coarse inference”. Forthcoming in *American Economic Journal: Microeconomics*.
- [30] Jehiel, P. (2005). “Analogy-Based Expectation Equilibrium”, *Journal of Economic Theory* 123: 81-104.
- [31] Kagel, J.H., and Roth, A.E. (1995). *The Handbook of Experimental Economics*. Princeton, New Jersey: Princeton University Press.
- [32] Manski, C. (2004). “Measuring expectations.” *Econometrica* 72, 1329-1376.
- [33] Marks, G. and Miller, N. (1987). “Ten years of research on the false-consensus effect: An empirical and theoretical review.” *Psychological Bulletin* 102: 72-90
- [34] McPherson, M., Smith-Lovin, L., and Cook, J. (2001). “Birds of a feather: homophily in social networks.” *Annual Review of Sociology* 27: 415-44.
- [35] Moody, A. (2013). “Adult anthropometric measures, overweight and obesity.” In Craig, R., and Mindell, J. *Health Survey for England 2012. Volume 1: Health, social care and lifestyles*. Health and Social Care Information Centre.
- [36] Moore, D.A., and Healy, P.J. (2008) “The trouble with overconfidence.” *Psychological Review* 115: 502-517.
- [37] Morris, S. (1995). “The common prior assumption in economic theory.” *Economics and Philosophy* 11: 227-253.

- [38] Loewenstein, G., O'Donoghue, T., and Rabin, M. (2003). "Projection Bias in Predicting Future Utility." *Quarterly Journal of Economics* 118: 1209-1248.
- [39] Park, A. and Sgroi, D. (2012). "Herding, contrarianism and delay in financial market trading" *European Economic Review* 56: 1020-1037.
- [40] Oswald, A. (2008). "On the curvature of the reporting function from objective reality to subjective feelings." *Economics Letters* 100: 369-372.
- [41] Robbins, H. (1952). "Some aspects of the sequential design of experiments." *Bulletin of the American Mathematical Society* 55, 527-535.
- [42] Ross, L., Greene, D., and House, P. (1977). "The 'false consensus effect': an egocentric bias in social Perception and Attribution Processes." *Journal of Experimental Social Psychology* 13: 279-301.
- [43] Rothschild, M. (1974). "A Two-Armed Bandit Theory of Market Pricing." *Journal of Economic Theory* 9: 185-202.
- [44] Rabin, M. (2002) "Inference by Believers in the Law of Small Numbers" *Quarterly Journal of Economics* 117(3): 775-816
- [45] Sgroi, D. (2003). "The right choice at the right time: A herding experiment in endogenous time" *Experimental Economics* 6, 159-180.
- [46] Tversky, A., and Kahneman, D. (1973). "Availability: a heuristic for judging frequency and probability." *Cognitive Psychology* 5, 207-232.
- [47] Tversky, A. and Kahneman, D. (1974). "Judgment under uncertainty: Heuristics and biases". *Science*, 185, 1124, 1128-1130.

Tables and Figures

Table 1: **Main Variables** Summary of the main characteristics of the 154 subjects participating the experiment: 84 males and 70 females. Note that there are 3 subjects that did not report their weight (1 male and 2 female) and 7 subjects who did not report their height (4 male and 3 female).

Variable	Mean	Std. Dev.	Min.	Max.	N
Weight Female	55.71	7.007	43	73	69
Weight Male	72.427	12.481	45	109	82
Height Female	162.418	7.237	138	181	67
Height Male	179.613	9.858	151	208	80
Political Stance	3.805	1.166	1	7	154
Happiness	4.766	1.021	1	7	154
Male	0.545	0.5	0	1	154

Table 2: **Determinants of Perceived Cumulative Distributions.** Dependent Variables: Perceived Share of Individuals below each individual Happiness (column 1), Political Stance (column 2), Weight (column 3) and Height (column 4). There are 154 subjects participating in the experiment. 3 subjects did not report their weight and 7 subjects did not report their height. Simple OLS estimator; Robust Standard Errors are reported in brackets.

	1 Happiness CDF b/se	2 Pol. Stance CDF b/se	3 Weight CDF b/se	4 Height CDF b/se
Happiness	0.6019*** (0.0427)			
Political Stance		0.4872*** (0.0506)		
Female*Weight			0.5502*** (0.0585)	
Male*Weight			0.5538*** (0.0682)	
Female*Height				0.7469*** (0.0627)
Male*Height				0.7907*** (0.0479)
Male	-0.0629** (0.0258)	0.0241 (0.0288)	0.0612 (0.0481)	0.1423*** (0.0394)
Constant	0.2630*** (0.0230)	0.2083*** (0.0278)	0.1327*** (0.0301)	0.0245 (0.0277)
r2	0.461	0.373	0.484	0.783
N	154	154	151	147

Table 3: **Determinants of Perceived Distributions.** Dependent Variables: Perceived Share of Individuals with the same mobile-phone brand (column 1) and the same information processing behavior, in terms of deferring to other individuals' information (column 2). There are 154 subjects participating in the experiment. 11 subjects reported that they did not own a mobile phone. Simple OLS estimator; Robust Standard Errors are reported in brackets.

	1	2
	Mobile Distribution	Deference Distribution
	b/se	b/se
Mobile	0.4288** (0.1825)	
Deference		0.4330*** (0.1060)
Male	-0.0557* (0.0289)	0.0141 (0.0324)
Constant	0.2251*** (0.0401)	0.3217*** (0.0394)
r2	0.059	0.097
N	141	154

Table 4: **Determinants of Perceived Cumulative Distributions of Happiness and Political Stance distinguishing Male and Female subjects.** Dependent Variables: Perceived Share of Individuals below each individual Happiness (column 1) and Political Stance (column 2). Simple OLS estimator; Robust Standard Errors are reported in brackets.

	1	2
	Happiness CDF	Pol. Stance CDF
	b/se	b/se
Female*Happiness	0.6137*** (0.0717)	
Male*Happiness	0.5957*** (0.0534)	
Female*Political Stance		0.3619*** (0.0766)
Male*Political Stance		0.5592*** (0.0633)
Male	-0.0577 (0.0386)	-0.0535 (0.0494)
Constant	0.2599*** (0.0284)	0.2542*** (0.0371)
r2	0.461	0.386
N	154	154

Figure 1: **Weight and Height:** The left panels represent the real CDFs (continuous lines) and the beliefs over the distributions (dots) and their respective Lowess function (dashed lines). The right panels represents the 45 degree lines (solid lines) and the local polynomial interpolation of the perceived and the real distributions (dashed lines), the shadow represents the 95% confidence interval.

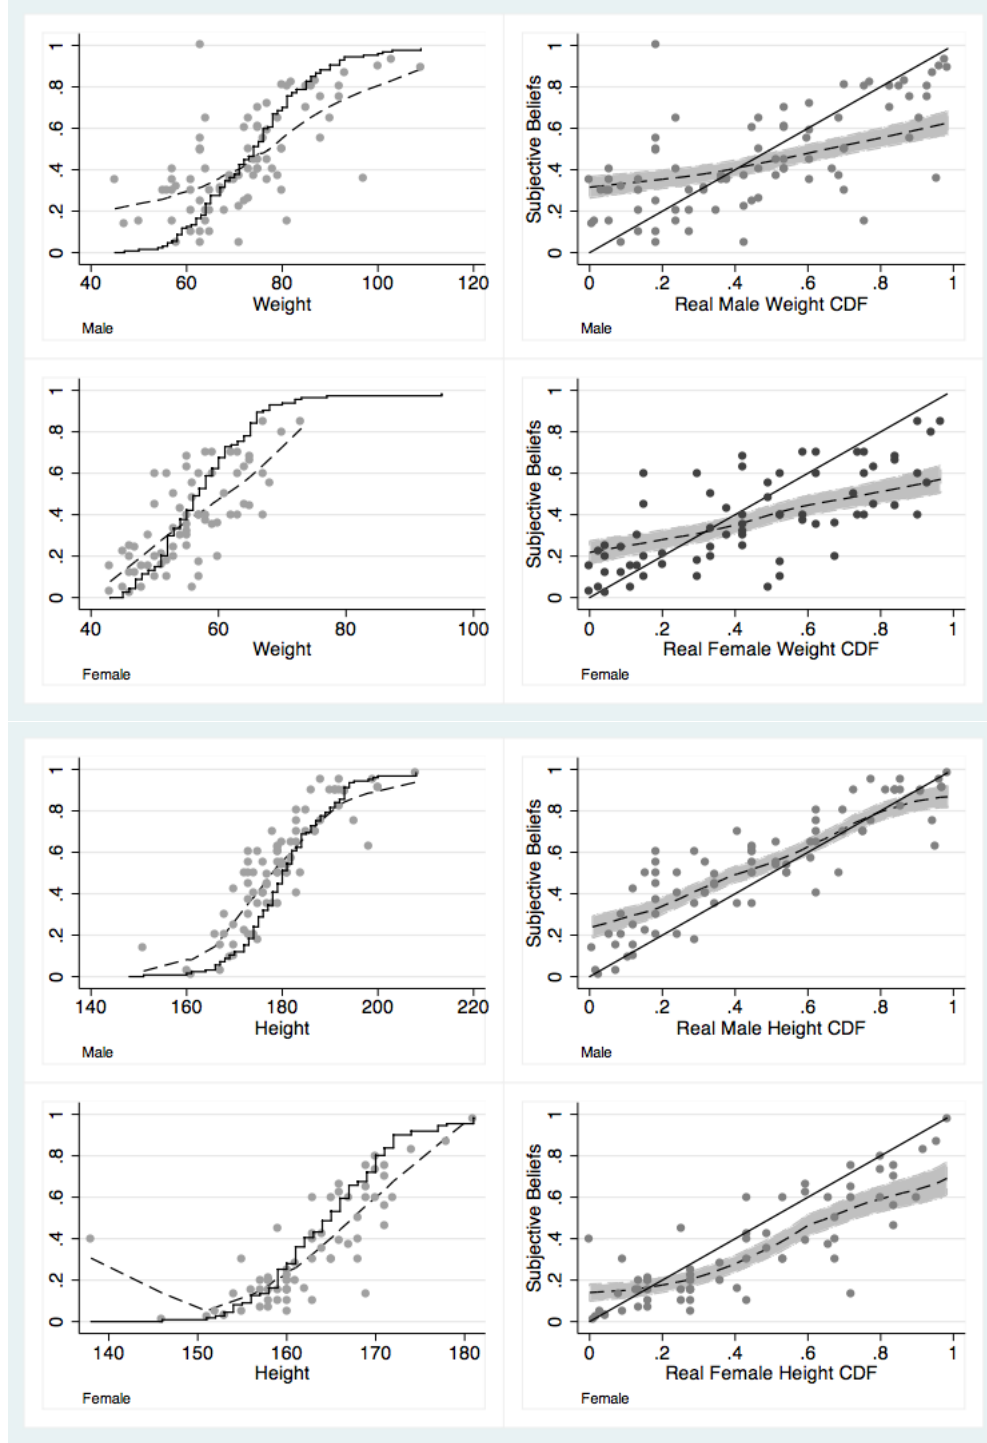


Figure 2: **Happiness and Political Stance:** The left panels represent the real CDFs (continuous lines) and the beliefs over the distributions (dots) and their respective Lowess function (dashed lines). The right panels represents the 45 degree lines (solid lines) and the local polynomial interpolation of the perceived and the real distributions (dashed lines), the shadow represent the 95% confidence interval.

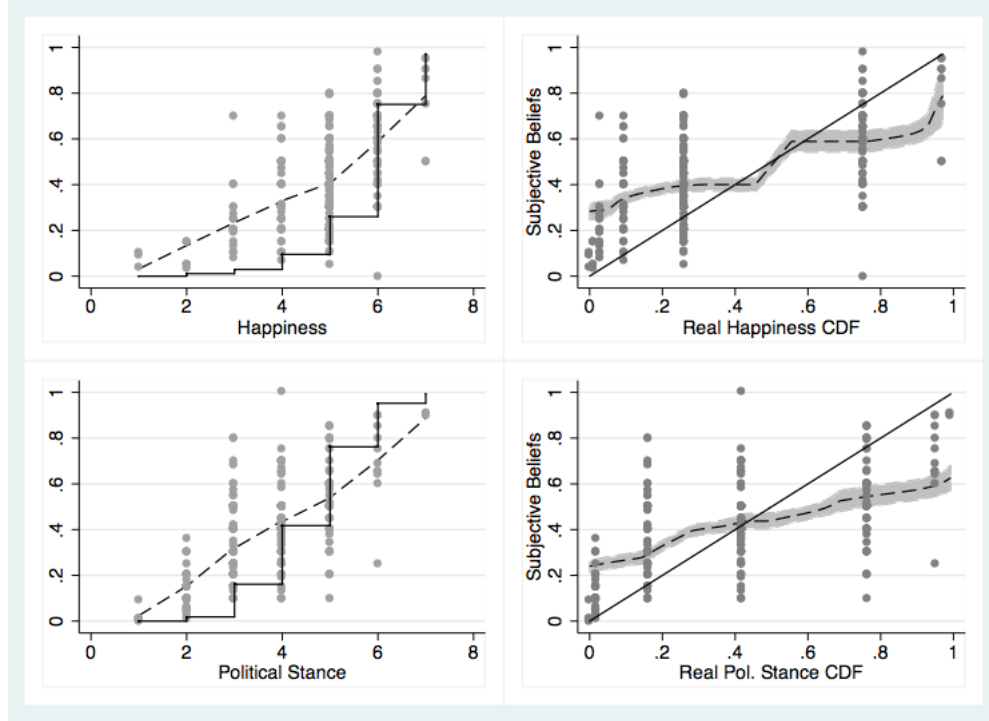


Figure 3: **Mobile Phones and Restaurant Choices:** The left panels represent the histograms of the simple real distributions (with the different characteristics ordered by frequencies), the beliefs over the distribution (dots) and their respective Lowess function (dashed lines). The right panels represents the 45 degree lines (solid lines) and the local polynomial interpolations of the perceived and the real simple distributions (dashed lines) with the shadow representing the 95% confidence interval.

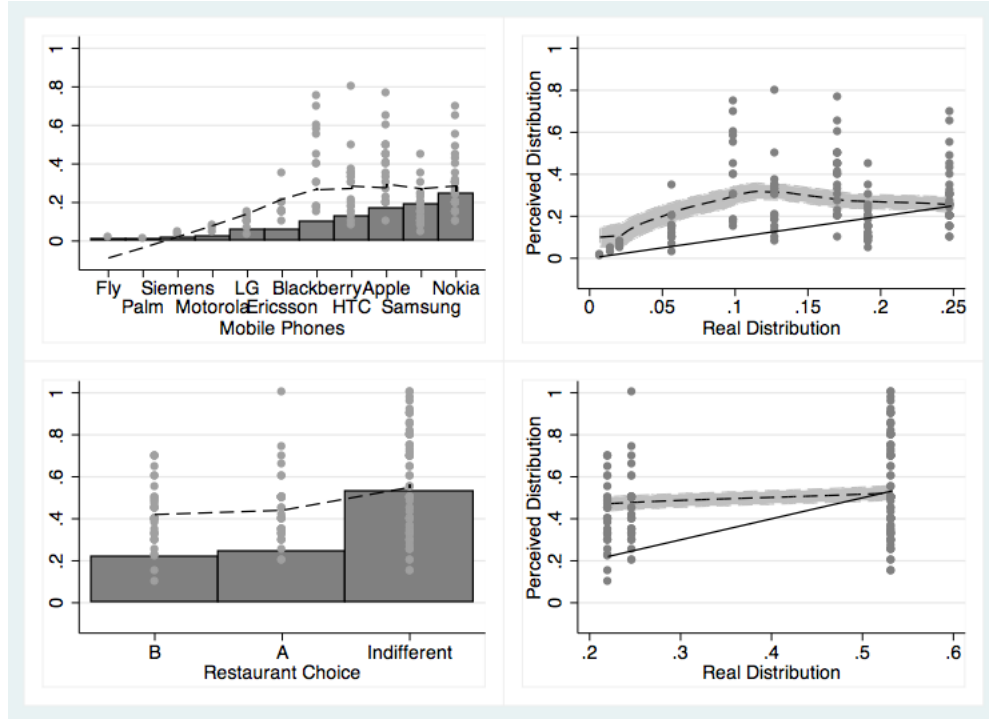


Figure 4: **Beliefs about averages of characteristics plotted against the individuals' own characteristics:** the solid lines represent the the linear interpolating function and the grey area represent the 95% confidence interval. The solid lines parallel to the axes represent the average values in our sample, in the y axis we also reported the confidence interval of the estimated averages.

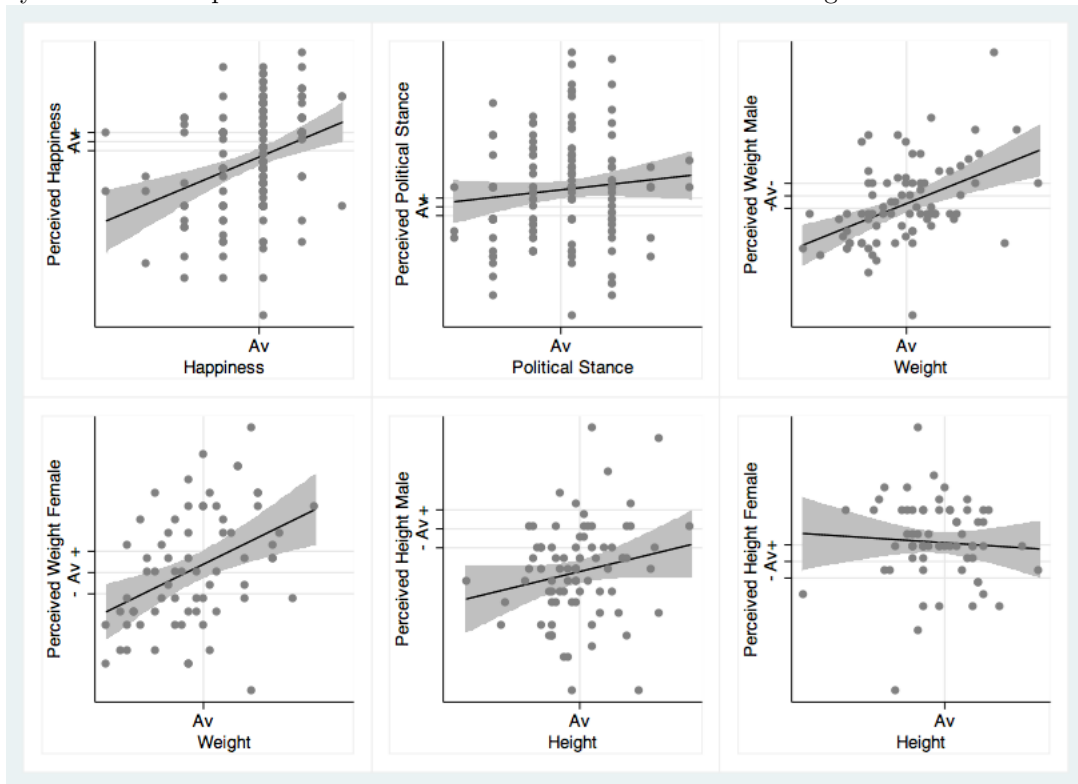


Figure 5: **Top, Bottom and Averages Beliefs:** For each characteristics we can observe the perceived averages of the different sections of the distributions and the real values in our sample. The bars represent the 95% confidence intervals.

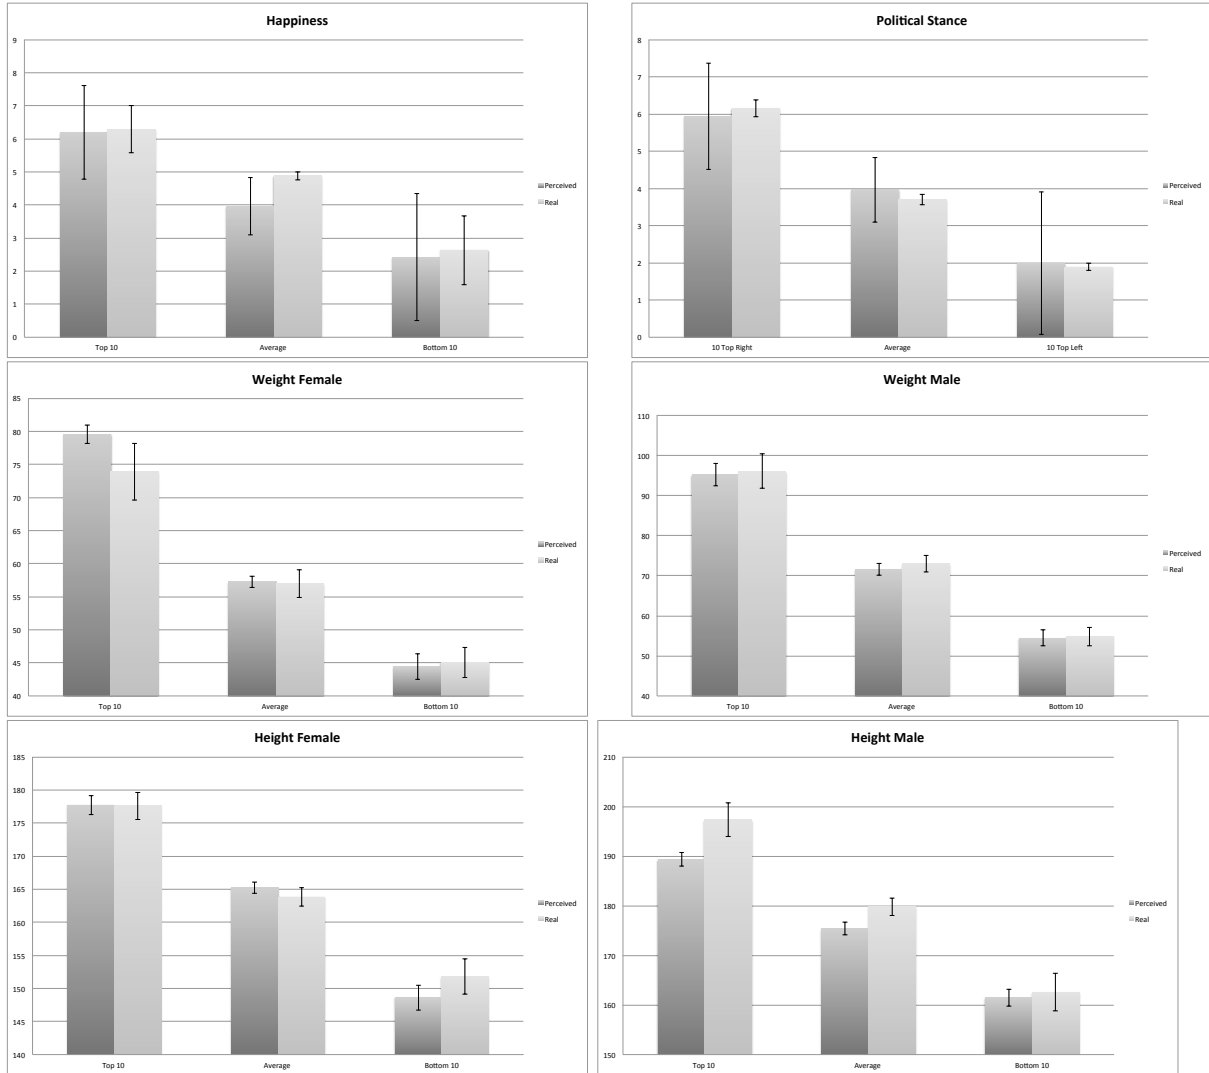


Table A.1: **Determinants of Perceived Cumulative Distributions without subjects belonging to the top and bottom 10% of each characteristic.** Dependent Variables: Perceived Share of Individuals below each individual Happiness (column 1) , Political Stance (column 2), Weight (column 3) and Height (column 4). Simple OLS estimator; Robust Standard Errors in Brackets.

	1 Happiness CDF b/se	2 Pol. Stance CDF b/se	3 Weight CDF b/se	4 Height CDF b/se
Happiness	0.5633*** (0.0457)			
Political Stance		0.4337*** (0.0566)		
Female*Weight			0.4893*** (0.0858)	
Male*Weight			0.5562*** (0.1108)	
Female*Height				0.7288*** (0.0746)
Male*Height				0.7785*** (0.0674)
Male	-0.0699** (0.0296)	0.0186 (0.0298)	0.0191 (0.0752)	0.1686*** (0.0486)
Constant	0.2777*** (0.0263)	0.2207*** (0.0313)	0.1597*** (0.0458)	0.0198 (0.0303)
r2	0.380	0.321	0.329	0.727
N	122	134	120	120

Table A.2: **Correlations between errors in estimating the different characteristics**

Variables	Happiness	Political Stance	Weight	Height	Mobile	Deference
Happiness	1.000					
Political Stance	0.078 (0.336)	1.000				
Weight	-0.055 (0.502)	-0.102 (0.215)	1.000			
Height	0.025 (0.762)	0.117 (0.159)	0.253 (0.002)	1.000		
Mobile	0.080 (0.343)	-0.036 (0.669)	0.033 (0.698)	0.047 (0.592)	1.000	
Deference	-0.016 (0.866)	0.072 (0.443)	0.132 (0.162)	-0.125 (0.192)	0.060 (0.546)	1.000