

CENTRE *for* ECONOMIC
P E R F O R M A N C E

CEP Discussion Paper No 1092

November 2011

Wishful Thinking

Guy Mayraz

Abstract

An experiment tested whether and in what circumstances people are more likely to believe an event simply because it makes them better off. Subjects observed a financial asset's historical price chart, and received both an accuracy bonus for predicting the price at some future point, and an unconditional award that was either increasing or decreasing in this price. Despite incentives for hedging, subjects gaining from high prices made significantly higher predictions than those gaining from low prices. The magnitude of the bias was smaller in charts with less subjective uncertainty, but was independent of the amount paid for accurate predictions.

JEL Classifications: D01, D03, D81, D84, G11

Keywords: Wishful-thinking, optimal expectations, priors and desires, payoff-dependent beliefs, asset prices

This paper was produced as part of the Centre's Wellbeing Programme. The Centre for Economic Performance is financed by the Economic and Social Research Council.

Acknowledgements

I am indebted to many colleagues for helpful comments and discussions, and especially to Michelle Belot, Gary Charness, Vincent Crawford, Erik Eyster, Paul Heidhues, Luis Miller, Matthew Rabin, Georg Weizsäcker, and Peyton Young. I gratefully acknowledge financial support from the Russell Sage Foundation Small Grants Program in Behavioral Economics and the John Fell OUP Research Fund.

Guy Mayraz is an Associate of the Wellbeing Programme, Centre for Economic Performance, London School of Economics.

Published by
Centre for Economic Performance
London School of Economics and Political Science
Houghton Street
London WC2A 2AE

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means without the prior permission in writing of the publisher nor be issued to the public or circulated in any form other than that in which it is published.

Requests for permission to reproduce any article or part of the Working Paper should be sent to the editor at the above address.

© G. Mayraz, submitted 2011

1 Introduction

People exhibit *wishful thinking* if they are more likely to believe an event simply because it makes them better off. The higher their utility if the event is true, the higher their subjective probability that it *is* true.

Wishful thinking has strong implications for decision making whenever subjective judgments of likelihood play an important role.¹ In single person decision problems agents may underestimate the risks they are exposed to, and overestimate the likelihood and value of uncertain rewards. In strategic environments the beliefs of agents with different interests may differ systematically in accordance with their interests. In dynamic environments a change in the payoff from an event may thereby alter the subjective likelihood that the event obtains, resulting in a change in preferences. For example, an investment in a financial asset would cause the investor to become more sanguine about the risks, leading her to escalate the investment well beyond her original plan.

These are, of course, all if-then statements: *if* wishful thinking affects beliefs in a given economic environment *then* it would have certain interesting implications for the behavior of decision makers in that environment. But is wishful thinking a real phenomenon? And supposing it is real, under what circumstances is it significant enough to affect decisions? What if the stakes are high?

The existing evidence for wishful thinking is mixed. On the one hand, wishful thinking provides a unified account for a wide range of empirical findings. On the other hand, while wishful thinking is the most parsimonious account of the evidence, each particular finding can also be given other explanations. Much of the evidence can be explained by cognitive biases that have nothing to do with wishful thinking (Miller and Ross, 1975; Nisbett and Ross, 1980). Other evidence can be explained by an ‘ego-utility’ bias, according to which people are more likely to believe an event if it fits their self-image (Kőszegi, 2006; Stone, 2006). The ego-utility account is related to wishful thinking, but is narrower in scope: whereas wishful thinking implies bias in beliefs over any payoff-relevant event, ego-utility implies bias only if the event reflects on the decision maker’s self-image. Finally, some of

¹As noted by Knight (1921) many decisions of interest to economists involve such judgments: “Business decisions, for example, deal with situations which are far too unique, generally speaking, for any sort of statistical tabulation to have any value for guidance. The conception of an objectively measurable probability or chance is simply inapplicable.” (III.VII.47); “Yet it is true, and the fact can hardly be overemphasized, that a judgment of probability is actually made in such cases.” (III.VII.40).

the findings that are suggestive of wishful thinking can also be given a plausible rational account (Van den Steen, 2004; Santos-Pinto and Sobel, 2005; Benoît and Dubra, 2011).

The relevant empirical findings appear under many different names, including self-serving bias, cognitive dissonance, over-optimism, and over-confidence. One powerful example is a study of self-serving bias in pre-trial bargaining, in which incentive compatible predictions for the amount the judge would award the plaintiff were significantly higher among subjects given the role of plaintiff than among subjects given the role of defendant (Loewenstein et al., 1993; Babcock et al., 1995; Babcock and Loewenstein, 1997), leading to an inefficient failure to reach a pre-trial settlement. Wishful thinking provides a natural explanation of the difference in beliefs between subjects in the two roles, since plaintiffs (defendants) are better-off if the award is high (low). It is also possible, however, to explain the difference in beliefs as the outcome of a cognitive bias to do with memory or attention. When preparing for the pre-trial bargaining phase of the experiment, subjects may have focused their reading of the evidence on arguments favoring their own side. A biased prediction of the judge's decision could then have resulted simply from basing the prediction on the part of the evidence that they best remembered.

The confound with ego utility is best exemplified by cognitive dissonance studies, such as the Knox and Inkster (1968) study of cognitive dissonance in horse betting. The key finding was that placing a bet on a horse significantly increases the confidence that the horse would win the race.² Since the act of placing the bet ties the bettor's utility to the horse's performance, the change in belief can be explained as the outcome of wishful thinking on the part of bettors. However, the horse's performance reflects on the bettor's judgment, as well as on her payoff more generally, and so the change in the beliefs of bettors is also consistent with an ego-utility explanation.

These and other empirical studies establish the existence of biased beliefs, and trace their consequences in economically important environments.³ These studies do not, however, make it possible to empirically separate wishful thinking from alternative explanations. The primary aim of the present paper is to offer a simple test of wishful thinking that is free of these confounds.

²By analogy, the acquisition of a financial asset would cause investors to become more positive about the asset.

³See also Olsen (1997), Camerer and Lovallo (1999), Malmendier and Tate (2005, 2008), Mullainathan and Washington (2009), Park and Santos-Pinto (2010), and Eil and Rao (2011).

The test uses a simple lab experiment. Subjects in the experiment observed a chart of historical wheat prices,⁴ and their one and only task was to predict what the price would be at some future time point. There was random assignment into two treatment groups: *Farmers*, whose payoff was increasing in the future price of wheat, and *Bakers*, whose payoff was decreasing in this price. Subjects in both groups also received a performance bonus as a function of the accuracy of their prediction.

Wishful thinking predicts bias whenever decision makers have a stake in what the state of the world is. *Farmers* gain from high prices, and their beliefs should therefore be biased upward as compared to what they would otherwise be. The opposite is true for *Bakers*. Given the random allocation, there should be a systematic difference in beliefs between the two groups, with *Farmers* expecting higher prices than *Bakers*.

In a standard rational model, there should be only random differences in beliefs between the two groups. The ego-utility explanation predicts bias if and only if the decision maker's self-image is involved. As this is not the case in the present experiment, ego-utility predicts no bias. Any number of cognitive biases could potentially affect beliefs in the experiment, and there is no good reason to expect predictions to coincide with the truth. However, since *Farmers* and *Bakers* are given the same information and the same prediction task, there is no plausible reason for cognitive biases to result in a difference in beliefs between the two groups. The presence or absence of a difference in beliefs between *Farmers* and *Bakers* therefore makes it possible to empirically separate wishful thinking from ego-utility, cognitive bias, and standard rational explanations.

The statistic used to identify a systematic difference in beliefs between the two groups was the difference between the average predictions of *Farmers* and *Bakers*. The prediction bonus formula was designed so that truthful reporting maximizes subjective expected payoff. As long as decision makers are risk-neutral over small amounts of money, the difference in predictions should provide an unbiased estimate of the difference in beliefs. Risk-averse subjects may, however, seek to intentionally hedge their predictions, so as to smooth their payoff across different states. Such hedging would result in *Farmers* under-reporting their true prediction, and an opposite bias for *Bakers*. Consequently, the estimated difference in beliefs between the two treatment groups may be biased downward.

The null hypothesis was defined as a non-positive difference in beliefs between *Farmers* and *Bakers*. Hedging could plausibly have resulted in a failure to reject

⁴Charts were adapted from real asset price data, though not specifically wheat prices.

the null when the true difference in beliefs is positive. There were no corresponding reasons to expect a false positive result. The actual observation was a positive and statistically measurable difference in predictions between *Farmers* and *Bakers* ($p < 0.0002$). This result is explained by wishful thinking, but by neither ego-utility nor cognitive biases.

The implication of this result is that wishful thinking does indeed affect subjective judgments of likelihood. However, while much of the interest in wishful thinking is in the possibility that it affects economically important decisions, the stakes in the experiment were relatively trivial amounts of money. The result would thus be considerably stronger if a similar experiment were run with much bigger prizes, but this may be prohibitively expensive. The alternative pursued in this paper is to test different theories of wishful thinking by varying the cost of getting beliefs wrong, while keeping the overall amounts small. Different theories of wishful thinking differ not only in the limit of high stake decisions, but also in their predictions for how the magnitude of the bias depend on the incentives for accuracy. If some particular theory survives this test, it can then be used with some confidence to make predictions for high stakes environments.

There are two principal approaches to modelling wishful thinking: strategic and non strategic. Psychologists commonly see wishful thinking as the outcome of a conflict between a desire to have an accurate view of the world and a desire to reach a ‘directional’ conclusion (Kunda, 1990). Economists mostly follow a similar approach, viewing wishful thinking as part of a strategy to balance the gain from positively biased beliefs against the cost in poor decisions by future selves (Akerlof and Dickens, 1982; Brunnermeier and Parker, 2005).⁵ From the present perspective, the most important implication of this family of models is that the magnitude of the bias decreases with the cost of getting beliefs wrong. We should thus expect an increase in the costs of biased beliefs to result in a measurably smaller bias, even if the stakes remain small.

Instead of modelling wishful thinking as a strategic choice, it is also possible to see it as the outcome of an imperfect belief formation process. Mayraz (2011) offers a descriptive model of wishful thinking along these lines, in which subjective beliefs are allowed to depend on the payoff consequences of events as well as on

⁵These are the two most relevant papers, but there are many other models in which agents choose an optimal level of bias in their beliefs, e.g. Carrillo and Mariotti (2000); Caplin and Leahy (2001); Benabou and Tirole (2002); Yariv (2002); Compte and Postlewaite (2004); Kőszegi (2006).

normatively relevant information. The key implication is that, controlling for the amount of available information, the magnitude of the wishful thinking bias should be independent of its effect on future decisions. In particular, a significant bias may well remain in high-stakes decisions, as long as the available information leaves sufficient room for uncertainty.

Differentiating between these two modelling approaches requires the ability to manipulate the incentives for holding accurate beliefs. The design of the experiment afforded a simple way to do so, by varying the scale of the accuracy bonus: the larger the potential bonus, the more subjects had to lose from holding biased beliefs. If wishful thinking is strategic, the magnitude of the bias should decrease in the scale of the accuracy bonus. If wishful thinking is non strategic, there should be no change in the magnitude of the bias as the scale of the accuracy bonus is increased.

Converting this intuition into a formal test requires quantitative predictions. ‘No change’ is a testable hypothesis, but ‘decreasing with the scale of accuracy bonus’ is not. Consequently, testing the hypothesis that wishful thinking is strategic made it necessary to focus on some particular strategic model. The best known such model is Optimal Expectations ([Brunnermeier and Parker, 2005](#)). Agents in this model have preferences over anticipated consumption, and choose beliefs in order to maximize their subjective expected utility. The constraint is that, once chosen, beliefs govern future choices and change only as the result of Bayesian updating. Agents therefore trade-off the gain from anticipating a high payoff, against the cost in a lower realized bonus: the more favorable they believe the future price to be, the higher is their anticipatory utility, but the lower the prediction bonus they can expect to receive. Increasing the scale of the accuracy bonus increases the cost of biased beliefs and reduces the optimal level of bias. Assuming risk-neutrality over small stakes, the quantitative prediction is that the magnitude of the bias would be inversely proportional to the scale of the accuracy bonus (Section 3.2).

Different sessions were run with different levels of accuracy bonus. The scale of the bonus was increased five fold, with the maximum bonus amount varying from £1 to £5. Results showed no decrease in the magnitude of the bias, consistent with the prediction of non strategic models. This result is statistically measurable: the prediction of the Optimal Expectations model was formally rejected ($p < 0.0140$), while that of non strategic models was not ($p < 0.4026$). The experiment, therefore, corroborates wishful thinking in its non strategic version. This, of course, is the version with the most far-reaching implications, implying that wishful thinking

affects any and all decisions based on subjective judgment, whatever the cost to the decision maker.

The Brunnermeier and Parker (2005) strategic model and the Mayraz (2011) non strategic model make comparative statics predictions not only for how the magnitude of the bias depends on the cost of holding biased beliefs, but also for its dependence on the amount of subjective uncertainty and on what subjects have at stake in the quantity that they form expectations over (the sensitivity of the final payoff to the day 100 price in the context of the present experiment). The prediction of both models is that the magnitude of the bias increases in both these factors. Testing these predictions cannot provide a further test of which model is correct, but it can provide some further assurance that the experiment is sensitive enough for the main conclusions to be trusted.

In order to make a test of the comparative statics of subjective uncertainty possible, subjects were asked to provide a confidence level together with their prediction. Confidence was provided on a 1-10 scale, calibrated with the help of examples provided as part of the instructions (Figure 3). By averaging the confidence reports across subjects, it was possible to obtain an estimate of the amount of subjective uncertainty in different charts. This made it possible to test the prediction that the bias in high subjective uncertainty charts is greater. Results were consistent with this prediction (Figure 5), and the null hypothesis that the magnitude of the bias is at least as high in low subjective uncertainty charts was rejected ($p < 0.0142$). A robustness test using a different measure of uncertainty yielded comparable results.

Due to insufficient data, a test of the comparative statics of the stakes was inconclusive. Two sessions were run with half the stakes, and the estimated bias was roughly half what it was in the baseline sessions. However, the null hypothesis that that the bias is the same could not be rejected.⁶

One concern with interpreting the results of the experiment is that subjects may have felt the task of predicting the day 100 price is impossible, and that they may as well choose whichever number they want to be true. Since *Farmers* gain from high prices and *Bakers* gain from low prices, *Farmers* would choose high guesses, and *Bakers* would choose low ones. If this explanation is correct, we would expect subjects who are generally confident in their predictions to be less biased than less

⁶The comparative statics of the stakes have been studied before in a different but related context. In a study of self-deception Mijović-Prelec and Prelec (2010) found a larger bias when stakes were higher (the ‘Anticipation Bonus’ treatment) as compared with lower bonus (the ‘Classification Bonus’ treatment).

confident subjects. Similarly, we would expect subjects who generally believe prices in financial markets are predictable to be less biased than subjects who do not think prices can be predicted. I tested the first prediction by defining a subject’s confidence level by the average confidence rating in her predictions across all charts. I tested the second prediction by asking subjects in the post experiment questionnaire whether they believe that prices in financial markets are generally predictable. In both cases I obtained just the opposite result: subjects who believe prices are predictable and relatively confident subjects are *more* biased than those who are less confident. These results suggest that this concern is misplaced. Moreover, they support the view that over-confidence is a manifestation of wishful thinking, and that the degree of wishful thinking bias is a stable individual characteristic.⁷

The remainder of the paper is organized as follows. Section 2 describes the experiment in detail. Section 3 develops the predictions of the Optimal Expectations (Brunnermeier and Parker, 2005) and Priors and Desires (Mayraz, 2011) models. Section 4 describes how the data were analyzed. Section 5 presents the results, and Section 6 concludes.

2 Experimental design

This section describes the experimental design. The implementation and protocol are in Section 2.1, and the specifics of the belief elicitation procedure in Section 2.2.

2.1 Implementation and protocol

The experiment was conducted at the Centre for Experimental Social Science (CESS) at Nuffield College, University of Oxford. The subject pool consisted of Oxford students who registered on the CESS website for participation in experiments. Business, finance, and economics students were excluded. A week before each session students meeting the sample restrictions received an email inviting them to participate in an experiment that would require one hour of their time. Further details were given on-site prior to the experiment itself. Registration was via an online form, allowing students to select one of several sessions, up to an upper limit of 14 students per

⁷For example, in Mayraz (2011) each person is characterized by a coefficient of relative optimism, which is a real-number characterizing her degree of optimism or pessimism. The same coefficient is assumed to determine her bias in all domains of subjective judgment.

session. Taking no-shows into account, sessions consisted of between 10 and 13 students. Altogether, 145 students took part in the experiment, of whom 57 were male and 88 female. The median age was 22.

Sessions were conducted in the afternoon over a total of six days. There were 12 sessions altogether. Half the sessions consisted of *Farmers*, and half of *Bakers*. The order of sessions was randomized in order to prevent any consistent relationship between the time of day in which a session was held, and the role given to the subjects who took part in that session.

After subjects were seated, they were each given a copy of the instructions, which they were able to refer to until the experiment ended. The instructions were also read aloud, and there was an opportunity for subjects to ask questions. The experiment itself consisted of 13 periods, the first of which was a training period, and the remaining 12 were earning periods. A given set of 13 charts was used throughout the experiment. One of these 13 charts was reserved for the training period, and the other 12 charts were used for the earning periods (Figure 2). The order of presentation was randomized independently between subjects. At the end of the experiment, each subject had one earning period chosen at random, and was paid in accordance with the payoff in that period.

The experiment was conducted in a computer lab, and was programmed using z-Tree (Fischbacher, 2007). Figure 1 shows an example of the interface. In each period subjects were shown a chart of wheat prices, and were asked to predict the price of wheat at some future date. Subjects were thus put in a somewhat similar position to speculators who ignore fundamental information, and predict future asset prices on the basis of historical price charts.⁸ In order to maximize the realism of the task, prices were adapted from real financial markets. The specific source was historical stock prices, scaled and shifted to fit into a uniform range. Charts were selected to include a variety of situations. Time was standardized across charts, so that all charts had space for prices going from day 0 to day 100. Subjects were only shown prices up to an earlier date, and the task was to predict what the price of wheat would be at day 100. The price range was also standardized, so that prices were always between £4,000 and £16,000.

After submitting their prediction, subjects were presented with a waiting screen until all other subjects had also made their prediction. There was therefore little

⁸Traders refer to the use of historical price charts in making buy and sell decisions as *Technical Analysis* (Murphy, 1999; Edwards and Magee, 2010).

or no incentive for speed. The transition to the next period only occurred after all the subjects in the room had submitted their prediction. A brief questionnaire was administered following the final period of the experiment. After all subjects completed the questionnaire, subjects were informed of their earnings, and were called to receive their payment.

Farmers were instructed that the price of wheat varies between £4,000 and £16,000, that it had cost them £4,000 to grow the wheat, and that they would be selling their wheat for the price that would obtain at day 100. Their notional profit was therefore between zero and £12,000, depending on the day 100 price. The payoff at the end of the experiment consisted of three parts: an unconditional £4 participation fee, profit from the sale of the wheat, and a prediction accuracy bonus. In the baseline sessions subjects received £1 in real money for each £1,000 of notional profit, and could earn up to an extra £1 from making a good prediction. The prediction procedure and bonus formula are explained in detail in Section 2.2. *Bakers* were told that they make bread, which they would sell for a known price of £16,000, and that in order to make the bread they would be buying wheat at the price that would obtain at day 100. The range of notion profit was therefore the same as that of *Farmers*, and all other particulars were also the same. The one difference was that that *Farmers* gained from high wheat prices, whereas *Bakers* gained from low prices.

Sessions differed in the scale of the accuracy bonus and in the stakes (the degree to which payoff depended on the price level at day 100). In the baseline sessions the maximum obtainable bonus was £1, and the amount received for each £1,000 of notional profit was also £1. Sessions were also conducted with a bonus level of £2 and £5, and with stakes of 50 pence for each £1,000 of notional profit.⁹ Table 1 lists the number of sessions in each condition.

2.2 The belief elicitation procedure

The belief elicitation procedure was designed with two goals in mind. The first was to make it possible to test for the presence or absence of wishful thinking bias, namely a systematic difference in beliefs between *Farmers* and *Bakers*. The second was to obtain a measure of the degree of subjective uncertainty in the predictions

⁹In sessions with lower stakes, subjects received an additional £3, so that the average payoff was the same as in the baseline sessions.

Table 1: The number of sessions for each combination of bonus scale and stakes.

bonus ^a	stakes ^b	sessions ^c	subjects
1	1	4	49 (25 <i>Farmers</i> , 24 <i>Bakers</i>)
2	1	2	26 (13 <i>Farmers</i> , 13 <i>Bakers</i>)
5	1	4	44 (23 <i>Farmers</i> , 21 <i>Bakers</i>)
1	0.5	2	26 (12 <i>Farmers</i> , 14 <i>Bakers</i>)

^a The amount in pounds subjects received for an optimal prediction of the day 100 price. The larger it was, the more subjects had to gain from holding accurate beliefs. The bonus for less good predictions was scaled accordingly.

^b The amount in pounds subjects received for each £1,000 of notional profit. The larger the stakes, the more subjects had to gain from the the day 100 price being high (if they were *Farmers*), or low (if they were *Bakers*).

^c Sessions were conducted in pairs: one for *Farmers* and the other for *Bakers*.

subjects make. This was important both for testing whether the magnitude of the bias is greater in charts with more subjective uncertainty, and for testing whether more confident individuals are also more biased.

In each period subjects were asked to report two numbers: a *prediction* and a *confidence level*. The prediction represented the expected day 100 price, and could be any number in the range of possible prices. The confidence level represented the (inverse of) the uncertainty in the prediction, and was reported on a 1-10 scale.

In order to give meaning to the 1-10 confidence scale, the instructions included visual examples of distributions with different prediction and confidence levels (Figure 3). The distributions were the weighted average of a normal distribution and a uniform one, with almost all the weight given to the normal. The prediction corresponded to the mean of the normal distribution, and the confidence level was inversely proportional to its standard deviation. The density corresponding to a prediction of $m \in [4000, 16000]$ and confidence level $r \in [1, 10]$ was

$$q(x) = (1 - \epsilon)\mathcal{N}(x|m, (\lambda r)^{-2}) + \epsilon \quad (1)$$

where $\mathcal{N}(\cdot|\mu, \sigma^2)$ represents a normal distribution with a given mean and variance, λ is a scale parameter, translating the 1-10 confidence scale into the scale of prices, and ϵ is the weight given to the uniform component. The effect of the latter was to ensure that the density was bounded below by ϵ , including at prices far from the prediction.

The scoring rule was logarithmic: subjects whose prediction and confidence level corresponded to a density q received a bonus given by

$$b(x) = \alpha \log (q(x)/\epsilon) \quad (2)$$

where x is the true day 100 price, and α is a parameter which determines the maximum bonus level.¹⁰ As $q \geq \epsilon$ (Equation 1), the bonus was positive for all possible predictions. The value of α was calibrated for the maximum bonus level in the session (Table 1).

To see under what conditions the scoring rule is incentive compatible, let P denote the probability measure representing the subject’s true beliefs, and suppose the subject reports a prediction m and a confidence level r . The subjective expectation of the bonus is given by the following expression:

$$\begin{aligned} \mathbb{E}_P[b(x)] &= \int p(x) \alpha \log \frac{q(x)}{\epsilon} dx = \alpha \left(\int p(x) \log \frac{q(x)}{p(x)} dx \right. \\ &\left. + \int p(x) \log p(x) dx - \log \epsilon \right) = \alpha \left(-D_{\text{KL}}(P||Q) - H(P) - \log \epsilon \right) \end{aligned} \quad (3)$$

where $D_{\text{KL}}(P||Q)$ is the *Kullback-Leibler divergence* (KL-divergence or relative entropy) between P and Q , and $H(P)$ is the entropy of P . Maximizing the expected bonus with respect to Q is thus equivalent to minimizing the KL-divergence $D_{\text{KL}}(P||Q)$. According to a standard result, $D_{\text{KL}}(P||Q) \geq 0$ for all Q , and is minimized if $Q = P$.¹¹

The scoring rule works best if subjects are risk neutral and beliefs are well approximated by a density in the family described by Equation 1. The scoring rule should then successfully elicit the prediction and confidence level for each subject in each chart, making it possible to identify the difference in beliefs between *Farmers* and *Bakers*, the average subjective uncertainty in each chart, and the average confidence for each subject.

One potential difficulty is hedging.¹² Consider a risk-averse *Farmer*. Her profit

¹⁰The logarithmic scoring rule was introduced in Good (1952). See Gneiting and Raftery (2007) for a recent discussion and comparison to other scoring rules.

¹¹This result, known as Gibb’s Inequality, follows directly from the fact that $\log x$ is a concave function (Cover and Thomas, 1991). The instructions explained that the expected bonus is maximized by reporting a prediction and confidence level that reflect the subject’s beliefs about the day 100 price. The bonus formula itself was included in a footnote.

¹²Blanco et al. (2008) find evidence of hedging in belief reporting when opportunities are trans-

is increasing in the price, and she would therefore prefer to receive the bonus in states in which the price is relatively low. Consequently, she could increase her subjective expected utility by reporting a lower number than her true beliefs. By a similar logic, a risk-averse *Baker* would be better-off by reporting a higher number. The result would be a downward bias in the estimated difference in beliefs between *Farmers* and *Bakers*.

A second potential problem is the possibility that the beliefs of some subjects are bi-modal, or otherwise not well approximated by a density in the family described by Equation 1. This could make it harder for subjects to see what prediction would maximize their payoff, making predictions within each group more variable than they would be otherwise. This increase in variance would translate into more noise in the estimated difference in beliefs between the two groups, though it should not result in bias.

3 Theory

This section formally develops the relevant predictions of different models of wishful thinking, including the standard model, Optimal Expectations (Brunnermeier and Parker, 2005), and Priors and Desires (Mayraz, 2011). By the standard model I mean the following two assumptions (i) choices maximize subjective expected utility, and (ii) subjective beliefs are independent of what a person has to gain or lose from an event being true. Optimal Expectations and Priors and Desires represent wishful thinking in its strategic and non strategic variants respectively. The ego-utility and cognitive bias explanations receive no separate treatment, as in the context of this experiment they agree with the assumptions of the standard model.

In developing the predictions of these three models I use the following timing framework: at $t = 0$ subjects observe a price chart and form their beliefs over the day 100 price; at $t = 1$ they report their prediction and confidence level, and consume anticipatory utility; at $t = 2$ the day 100 price is revealed, and payoffs are realized. In order to simplify the analysis, I assume that subjects are risk neutral and that their beliefs about the day 100 price can be represented by a distribution from the family described by Equation 1. Given these assumptions, the prediction made at $t = 1$ coincides with the $t = 1$ beliefs.

parent and incentives are strong. [Armantier and Treich \(2010\)](#) discuss hedging in probability elicitation.

3.1 The standard model

Different individuals in the standard model may end up with different beliefs for unmodelled (random) reasons. By assumption, however, a person's subjective beliefs about the day 100 price are not affected by whether she is assigned the role of *Farmer* or *Baker*. Since the role allocation is random, it follows that the $t = 0$ subjective beliefs of *Farmers* and *Bakers* are drawn from the same distribution. Since the prediction coincides with the $t = 1$ beliefs, and as no new information is observed between $t = 0$ and $t = 1$, it follows that predictions are also drawn from the same distribution. Consequently, there is no systematic difference in predictions between *Farmers* and *Bakers*.

3.2 Optimal Expectations

Optimal expectations agents choose their prior beliefs in order to maximize their discounted subjective expected utility, where each period's instantaneous utility includes anticipatory utility as well as standard consumption utility.

The payoff in the experiment is realized at $t = 2$, and consists of two components: the profit and the accuracy bonus. The profit is a function of the true price, while the bonus depends on the accuracy of the $t = 1$ beliefs. Anticipatory utility is proportional to the expected value of the profit and bonus, with expectations computed using the $t = 1$ beliefs. The more optimistic those beliefs are, the higher is anticipatory utility, but the less accurate the prediction is likely to prove. The $t = 0$ decision maker choosing her $t = 1$ beliefs therefore faces a trade-off: more bias increases the anticipatory utility experienced at $t = 1$, but lowers the expected value of the $t = 2$ consumption utility.

Let P and Q denote the probability distributions representing the $t = 0$ and $t = 1$ beliefs respectively. At $t = 0$ the agent maximizes a weighted sum of the $t = 1$ anticipatory utility and $t = 2$ realized payoff. Let η denote the weight given to anticipatory utility, so that the weight given to the realized payoff is $1 - \eta$. Letting x denote the true day 100 price, the profit can be written as $\phi\kappa x + l$, where x is true day 100 price, κ represents the stakes (the absolute value of the slope relating the profit to the day 100 price), and ϕ denotes the direction, with $\phi = 1$ for *Farmers* and $\phi = -1$ for *Bakers*. I denote the bonus by $b(x)$, where b is defined by Equation 2.

The $t = 0$ maximand can thus be written as follows:

$$W = \eta \mathbb{E}_Q[\phi \kappa x + b(x)] + (1 - \eta) \mathbb{E}_P[\phi \kappa x + b(x)] + l \quad (4)$$

In order to derive the comparative statics of the bias in closed form I make a couple of simplifying assumptions. First, I assume that P and Q are normal: $P = \mathcal{N}(\mu_0, \sigma_0^2)$, and $Q = \mathcal{N}(\mu_1, \sigma_1^2)$. Second, I assume that only the mean of Q is subject to bias, i.e. $\sigma_1 = \sigma_0 = \sigma$. Given these assumptions and using Equation 3, we can rewrite Equation 4 as follows:

$$\begin{aligned} W &= \eta \mathbb{E}_Q[\phi \kappa x + b(x)] + (1 - \eta) \mathbb{E}_P[\phi \kappa x + b(x)] + l \\ &= \eta(\phi \kappa \mu_1 - \alpha H(Q) - \alpha D_{\text{KL}}(Q||Q) - \alpha \log \epsilon) \\ &\quad + (1 - \eta)(\phi \kappa \mu_0 - \alpha D_{\text{KL}}(P||Q) - \alpha H(P) - \alpha \log \epsilon) + l \\ &= \eta(\phi \kappa \mu_1 - \alpha H(Q)) - (1 - \eta)\alpha D_{\text{KL}}(P||Q) + C \end{aligned} \quad (5)$$

where C collects factors that are independent of Q . The two terms that depend on Q represent, respectively, the *gain* in anticipatory utility from adopting optimistic beliefs, and the *cost* in expected realized payoff of adopting such beliefs. The gain term has two components. The first represents the anticipated profit, and is proportional to $\mu_1 = \mathbb{E}_Q[x]$. The second represents the anticipated bonus, and is decreasing in the degree of uncertainty in Q , measured by its entropy $H(Q)$. The gain term is thus larger the more favorable is the expected day 100 price, and the more certain the subject is about her prediction. The cost term represents the reduction in expected bonus due to the bias in the prediction that follows from the bias in the $t = 1$ beliefs, and is proportional to the Kullback-Leibler divergence between the $t = 0$ beliefs P and the $t = 1$ beliefs Q . Thus, if the subject cared only about the realized payoff she would choose not to bias her beliefs at all ($Q = P$). If, instead, she cared only about her $t = 1$ instantaneous utility, she would choose to believe that the most favorable price would be realized,¹³ and would further choose to assign this prediction as little subjective uncertainty as possible.

If η is sufficiently small, the optimal choice of μ_1 would be an extreme value in the favorable direction. Otherwise, the optimal value of μ_1 would be at an internal point, where $\partial W / \partial \mu_1 = 0$. Since we do not observe subjects making extreme

¹³That is, the highest possible price of £16,000 if a *Farmer*, and the lowest possible price of £4,000 if a *Baker*.

predictions I assume that η is large enough that the optimal value of μ_1 is at an internal point. Using the standard formula for the KL-divergence between two normal distributions (Johnson and Sinanovic, 2001), and noting that $H(Q)$ is independent of μ_1 , the derivative can be written as follows:

$$\begin{aligned}\frac{\partial W}{\partial \mu_1} &= \eta\phi\kappa + \eta\frac{\partial H(Q)}{\partial \mu_1} - (1-\eta)\alpha\frac{\partial D_{\text{KL}}(P\|Q)}{\partial \mu_1} \\ &= \eta\phi\kappa - (1-\eta)\alpha\frac{(\mu_1 - \mu_0)}{\sigma^2}\end{aligned}\tag{6}$$

Setting the derivative to zero and solving for μ_1 we obtain the following expression for the bias:

$$\mu_1 - \mu_0 = \phi\left(\frac{\eta}{1-\eta}\right)\left(\frac{\kappa\sigma^2}{\alpha}\right)\tag{7}$$

where κ represents the stakes, or the degree to which the profit is dependent on the value of the day 100 price, σ^2 represents the degree of subjective uncertainty, and α represents the scale of the accuracy bonus, or the cost of holding biased beliefs.

Equation 7 describes the bias in the beliefs of one particular individual. The prediction for the average bias in the population of subjects in the same role is

$$\mathbb{E}[\mu_1 - \mu_0] = \mathbb{E}[\mu_1] - \mathbb{E}[\mu_0] = \phi\mathbb{E}\left[\frac{\eta}{1-\eta}\right]\left(\frac{\kappa\sigma^2}{\alpha}\right)\tag{8}$$

where I allow for the possibility that η varies between individuals, but assume that it is independent of σ^2 (because of the random assignment η is independent of κ and α). Finally, it also follows from the random allocation that the undistorted beliefs of *Farmers* and *Bakers* are drawn from the same distribution, and that in particular $\mathbb{E}\mu_0$ is the same in both groups. The expected difference in beliefs between the two groups is therefore given by

$$b_{\text{optimal expectations}} = 2\mathbb{E}\left[\frac{\eta}{1-\eta}\right]\left(\frac{\kappa\sigma^2}{\alpha}\right) \propto \frac{\kappa\sigma^2}{\alpha}\tag{9}$$

Optimal Expectations thus implies a systematic difference in beliefs between *Farmers* and *Bakers* that is proportional to the stakes and to the degree of subjective uncertainty, and inversely proportional to the cost of getting beliefs wrong.

3.3 Priors and Desires

Priors and Desires (Mayraz, 2011) is a non strategic model, that starts by allowing for the possibility that a person’s beliefs may depend on what she wants to be true. The latter is formalized by a *payoff-function*, which is a mapping from states to utility values, representing the dependence of the decision maker’s utility on the state of the world. A number of simplifying assumptions are made, and a representation is derived. The bias in the subjective beliefs of a person with a payoff-function f is represented by the following equation:

$$q(s) \propto p(s)e^{\psi f(s)} \quad (10)$$

where s denotes the state, q represents the decision maker’s actual (distorted) beliefs, and p represents her undistorted beliefs, or the beliefs she would hold if she were indifferent between all states.¹⁴ Finally, ψ is a real-valued parameter, called the *coefficient of relative-optimism*, which describes how optimistic or pessimistic that particular individual is. A positive value corresponds to optimism, a negative value to pessimism, and a zero value to realism.

The payoff-function in the experiment is the mapping linking the subject’s payoff to the day 100 price.¹⁵ Using the same notation as in Section 3.2, the payoff-function is given by $f(x) = \phi\kappa x + l$, where x is the day 100 price, κ represents the stakes, or the slope relating payoff to the day 100 price, and ϕ denotes the direction, with $\phi = 1$ for *Farmers* and $\phi = -1$ for *Bakers*. Suppose, as in Section 3.2, that undistorted beliefs are given by a normal distribution $P = \mathcal{N}(\mu_0, \sigma^2)$. The prediction for the actual (distorted) beliefs can be obtained using Equation 10. Mayraz (2011, proposition 3) analyzes the case of a normal distribution with a linear payoff-function, and shows that the distorted beliefs are normally distributed with the same variance, and that the mean is shifted in proportion to the coefficient of relative optimism ψ , the stakes, and the variance. In other words, the distorted probability measure is given by $Q = \mathcal{N}(\mu_1, \sigma^2)$, where

$$\mu_1 - \mu_0 = \phi\psi\kappa\sigma^2 \quad (11)$$

This equation describes the bias in the beliefs of some particular individual, and is

¹⁴That is, if $f(s) = f(s')$ for all s and s' .

¹⁵In principle, it should be the payoff in utility terms, but I am assuming throughout this section that subjects are risk neutral over small amounts of money.

the Priors and Desires analogue of Equation 7. By analogy with Section 3.2, the expected difference in beliefs between *Farmers* and *Bakers* is

$$b_{\text{priors and desires}} = 2\mathbb{E}[\psi]\kappa\sigma^2 \propto \kappa\sigma^2 \quad (12)$$

Comparing this result to Equation 9, we see that—as with Optimal Expectations—the magnitude of the bias is proportional to the stakes κ and the degree of subjective uncertainty σ^2 . However, whereas in Optimal Expectations the magnitude of the bias is inversely proportional to the cost of getting beliefs wrong α , the magnitude of the bias in Equation 12 is independent of α .

4 Analysis

This section describes how the data was analyzed. The resulting estimates are presented in Section 5.

4.1 Minimizing hedging bias

As noted in Section 2.2, hedging could lead to a downward bias in estimating the difference in beliefs between *Farmers* and *Bakers*. In order to minimize this risk, a questionnaire was administered after the experiment itself was concluded, in which subjects were asked whether they always reported their best guesses, or whether they sometimes reported a higher or lower number. Out of a total of 145 students who took part in the experiment, 132 claimed to have always reported their best guess, and 13 admitted to an intentional bias in their predictions. Observations from these 13 subjects were excluded from the main analysis.

4.2 Main regression

The raw data from the experiment consist of the predictions and confidence levels reported by individual subjects in individual charts. The primary goal in analyzing the data was to determine whether predictions were affected by wishful thinking. Let $y_{i,j}$ denote the prediction made by subject i in chart j , and let $t_i \in \{1, -1\}$ denote whether subject i is a *Farmer* or a *Baker*. We want to know whether y_{ij} is systematically higher if $t_i = 1$. In order to answer this question formally I used the

following regression model:

$$y_{ij} = 0.5\beta t_i + \sum_j \gamma_j d_j + \epsilon_{ij} \quad (13)$$

where d_j is a dummy for chart j , and ϵ_{ij} is the error term. The value of β represents the contribution of wishful thinking. The null hypothesis is that $\beta \leq 0$.

4.3 Comparative statics

The second goal in analyzing the data was to investigate the comparative statics of the bias. This required estimating the bias separately in different subsamples of interest. Let K denote a partition of the sample, indexed by k , and let c_{ijk} denote a dummy for whether the prediction of subject i in chart j belongs to subsample k . Assuming wishful thinking is the only systematic source of difference in predictions between subjects, we can generalize Equation 13 as follows:

$$y_{ij} = 0.5 \sum_{k \in K} \beta_k c_{ijk} t_i + \sum_j \gamma_j d_j + \epsilon_{ij} \quad (14)$$

In this equation β_k represents the average difference in predictions between *Farmers* and *Bakers* in class k , and can be used to define formal comparative statics hypotheses.

4.4 Standard errors

Unobserved factors may result in a correlation in the predictions made by the same subject in different charts, so that ϵ_{ij} may be correlated with ϵ_{ik} for $j \neq k$. In order to allow for this possibility, standard errors are clustered by subject in all regressions.

5 Results

This section presents the results of the experiment, starting with the overall difference in predictions between *Farmers* and *Bakers*, and continuing with the comparative statics of the bias. Parameter estimates and statistical test results are presented

in summary form in Table 3. Figures 4 and 5 provide a graphical illustration of the results.

5.1 Wishful thinking bias

The overall magnitude of the wishful thinking bias corresponds to the systematic difference in predictions between *Farmers* and *Bakers* across the entire sample, represented by the value of β in Equation 13. The estimate for this number is £452, measured with a robust standard error of £123. The null-hypothesis that it is non-positive is rejected with a p -value of 0.0002.

This estimate excludes observations from the 13 subjects who admitted in the post experiment questionnaire to biased reporting of their beliefs (Section 4). If these subjects are nonetheless included, the estimate goes down to £390. This difference is consistent with the prediction that risk-averse *Farmers* (*Bakers*) would intentionally understate (overstate) their estimates of the day 100 price.

The observed difference in predictions between *Farmers* and *Bakers* can be explained by wishful thinking, but not by ego-utility or by a cognitive bias.

5.2 Incentives for accuracy

Strategic models of wishful thinking predict that the magnitude of the bias would be decreasing in the incentives for accuracy, while non strategic models predict that it would remain the same. In order to determine whether higher incentives for accuracy result in lower bias, Equation 14 was used to estimate the difference in beliefs between *Farmers* and *Bakers* separately in sessions with different levels of accuracy bonus (Table 1).

The results in Table 3 are that the estimated bias is actually *greater* in sessions with a higher bonus level, the point estimates being 298, 560, and 645, respectively. On the face of it, these results are consistent with neither type of model. Formal testing, however, reveals that the apparent increase in the magnitude of the bias may well be random ($p < 0.4026$). The data is, therefore, consistent with the prediction of non strategic models that the magnitude of the bias would be invariant to changes in the incentives for holding accurate beliefs.

The same is not true, however, for strategic models. The prediction of the Optimal Expectations model is that the magnitude of the bias would be inversely proportional to the scale of the accuracy bonus (Section 3.2). That is, the bias in

£2 bonus sessions should be half the size of the bias in £1 bonus sessions, and the bias in £5 bonus sessions should be one fifth the size. This prediction is rejected by the data ($p < 0.0140$).¹⁶

The first panel of Figure 5 shows these results graphically. Though the point estimates are increasing in the maximum level of the accuracy bonus, a horizontal parallel line can be comfortably fitted within the confidence intervals. The same is not true, however, for a hyperbolic curve.

5.3 Subjective uncertainty

According to both Optimal Expectations (Section 3.2) and Priors and Desires (Section 3.3), the magnitude of the bias should be increasing in the degree of subjective uncertainty. In order to test this prediction, I divided the 12 charts used in the paying periods into two equal sized groups by the degree of subjective uncertainty in the chart, and used Equation 14 to estimate the bias separately in the two subsamples.¹⁷ I used two different measures of subjective uncertainty. The first was based on the confidence ratings that subjects provided: charts were classified into the high (low) subjective uncertainty group if the mean (across all subjects) of the confidence rating for the chart was below (above) median. The second measure of uncertainty was the within group variance of predictions: charts were classified into the high (low) subjective uncertainty group if the within group variance of predictions for that chart was above (below) median. In practice, the two measures resulted in nearly identical classifications.

Depending on the measure used, the estimated bias was 635 or 677 in the group of high subjective uncertainty charts, and 269 or 227 in the low subjective uncertainty group. The null hypothesis—that the magnitude of the bias in high subjective uncertainty charts would be less than or equal to the magnitude of the bias in low subjective uncertainty charts—was rejected with a p -value of 0.0142 when using the first classification method, and a p -value of 0.0034 when using the second (Table 3).

These results support the qualitative prediction that the magnitude of the bias is increasing in the degree of subjective uncertainty. Given that the qualitative prediction of the two models fits the data, it is interesting to try and test the specific functional form predicted by the two models. The quantitative prediction

¹⁶This is the a joint hypothesis test. The hypothesis that the magnitude of the bias in £5 sessions is one fifth that of £1 sessions is rejected with a p -value of 0.0069.

¹⁷Each subsample consists of observations from all subjects, but in only half the charts.

is that the magnitude of the bias is linear in the variance of subjective uncertainty. The following equation should thus prove to be a better model of the data than Equation 13:

$$y_{ij} = 0.5\beta'\sigma_j^2t_i + \sum_j \gamma_j d_j + \epsilon_{ij} \quad (15)$$

In this equation the $0.5\beta t_i$ term in Equation 13 is replaced by $0.5\beta'\sigma_j^2t_i$, where σ_j^2 is the variance of subjective uncertainty in chart j .

Testing this quantitative prediction requires a good proxy for the variance of subjective uncertainty. Using the above measures of subjective uncertainty, we can identify σ_j^2 either with the square of the inverse mean confidence rating in chart j , or with the mean within group prediction variance for chart j .¹⁸ Table 2 shows the resulting regression fit when estimating the two equations using both proxies for the variance of subjective uncertainty, as well the results of fitting a model which includes both the $0.5\beta t_i$ term of Equation 13 and the $0.5\beta'\sigma_j^2t_i$ of Equation 15. The results show that Equation 15 indeed provides a better fit to the data, consistent with the prediction that the magnitude of the bias is linear in the degree of subjective uncertainty.

The same results can also be seen graphically in the second and third panels of Figure 5. Panel 2 plots the estimated wishful thinking bias in the 12 charts against the mean prediction confidence in the chart, and panel 3 plots the same data against the within group prediction variance. In both panels a curve is fitted to the data using Equation 15.

5.4 Stakes

Optimal Expectations and Priors and Desires also predict that the magnitude of the bias is increasing in the stake subjects have in what the day 100 price would be. Payoff depends on the day 100 price via the notional profit, which is linear in the day 100 price with a slope of 1. The amount of money received for each £1,000 of notional profit was £1 in 10 sessions and 50p in the remaining 2 sessions (Table 1).

I estimated the magnitude of the bias separately in these two subsamples (Equation 14). The magnitude of the bias was 260 in the low stakes subsample, and 495 in the standard stakes subsample. These results are consistent with the prediction that the magnitude of the bias is linear in the stakes ($p < 0.9668$). The modest

¹⁸This assumes a representative agent approximation.

Table 2: Testing whether the magnitude of the bias increases with the variance of subjective uncertainty. Column 1 fits a model in which the bias is independent of subjective uncertainty (Equation 13). Column 2 and 4 fit a model in which the magnitude of the bias is linear in the variance (Equation 15). Columns 3 and 5 fit a model which allows for both regressors. Method 1 and method 2 refer to the two proxies for subjective uncertainty (Section 5.3). The $t_i\sigma_j^2$ variable is normalized to have the same standard deviation as t_i , so that the regression coefficients are comparable in size. Robust standard errors are in parentheses. The regression R^2 is computed after netting out the contribution of the chart dummies. Statistical significance indicators: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

	method 1		method 2	
t_i	452*** (122)	-473* (259)	-458* (272)	
$t_i\sigma_j^2$	497*** (129)	955*** (309)	503*** (130)	945*** (319)
R^2	0.0181	0.0218	0.0230	0.0237

variance in the stakes between sessions was, unfortunately, insufficient to produce statistically measurable results, and the hypothesis that the bias is not any smaller in the low stakes subsample could not be rejected ($p < 0.2313$). See also Table 3 and panel 4 of Figure 5.

5.5 Over-confidence

Section 5.1 demonstrates the existence of a systematic difference in predictions between *Farmers* and *Bakers*. This difference in predictions is interpreted as evidence of wishful thinking bias affecting subjects' judgment about the day 100 price. A key assumption is that subjects believe they have better than random odds of making a good prediction, so it is in their interest to report their true beliefs. If this assumption is not true, subjects could very well choose whichever prediction they enjoy making, without having to worry about losing the prediction bonus. As long as subjects prefer reporting a price that would benefit them, we could observe a systematic difference in predictions between *Farmers* and *Bakers* that has nothing to do with wishful thinking.

If this alternative explanation is correct, we would expect subjects who lack

confidence in their predictions to be more biased than confident subjects, since such subjects have less to lose from biasing their prediction. Similarly, we would expect subjects who generally believe prices in financial markets are unpredictable to be more biased than subjects who believe prices can be predicted.

In order to test the first prediction I defined a proxy for a subject's confidence by the average prediction confidence for that subject across all charts. I then split the sample into more and less confident subjects, and estimated the bias separately in the two subsamples. In order to test the second prediction I included a question in the post experiment questionnaire about the predictability of prices in financial markets, and divided subjects into two groups by whether they thought prices can generally be predicted. The bias was then estimated separately in the two subsamples.¹⁹

The result was just the opposite: subjects who believe prices are predictable and relatively confident subjects are *more* biased than those who are less confident. Specifically, the estimated bias among relatively confident subject is 628, compared with 276 among less confident subjects. The hypothesis that more confident subjects are less biased is rejected with p -value of 0.0732. Similarly, the estimated bias among subjects who believe prices in financial markets to be generally predictable was 613, as compared with 292 among subjects who believed prices cannot be predicted. The hypothesis that subjects who believe prices to be predictable are less biased was rejected with a p -value of 0.0997.

By and large, therefore, subjects believe they have at least some ability to predict the day 100 price, and the stronger this belief is, the *more* biased they are. This result is consistent with the wishful thinking interpretation, and further suggests that over-confidence is a manifestation of wishful thinking, and that the degree of wishful thinking bias is a stable individual characteristic.²⁰

¹⁹The question was “We are interested in what people believe about financial markets. How predictable are the movements of prices in financial markets in your opinion?” The possible choices were: “Prices can be predicted to a significant extent”, “Prices can rarely be predicted”, and “The idea that prices can be predicted is an illusion”. The first choice was defined as *yes*, and the other two as *no*. The distribution of answers was 66, 58, and 8, respectively.

²⁰This explains why individuals with more than average wishful thinking bias also tend to be over-confident. The tendency to be more or less biased can be identified with the coefficient of relative optimism in [Mayraz \(2011\)](#).

5.6 Gender

Though the psychology evidence is mixed (Lundeberg et al., 2000), certain behavioral differences between men and women, such as a propensity to overtrade among men, have been interpreted as evidence of gender differences in confidence (Barber and Odean, 2001). Subjects in the experiment included 62 percent females and 38 percent males, and there was therefore sufficient variation to test for gender differences in wishful thinking. The estimated bias is 411 for males and 477 for females, and the hypothesis of no difference cannot be rejected ($p < 0.7956$).

6 Conclusion

This paper describes an experimental test of wishful thinking bias in predictions of asset prices. Subjects received an accuracy bonus for their predictions of the future price of an asset, and an unconditional payment that was either increasing or decreasing in this price. Both groups of subjects had the same information, and faced the same incentives for accuracy. Nevertheless, and despite incentives for hedging, subjects in the group benefiting from high prices predicted systematically higher prices than subjects in the group benefiting from low prices. These results are consistent with wishful thinking, and cannot be accounted for by such alternative explanations as ego-utility or cognitive bias.

By varying the scale of the accuracy bonus it was possible to test whether the magnitude of the bias decreases with the incentives to hold accurate beliefs. No such decrease was found, and the prediction of Optimal Expectations (Brunnermeier and Parker, 2005) that the magnitude of the bias is inversely proportional to the incentives for accuracy, was formally rejected. This result is hard to square with strategic models of wishful thinking, but is consistent with non strategic models, such as Mayraz (2011). The implication is that wishful thinking can significantly affect beliefs even if the costs are high.

Other comparative statics results include good evidence that wishful thinking bias is stronger when subjective uncertainty is high, evidence that over-confidence and wishful thinking bias go together, and some evidence of greater bias when payoff is more strongly dependent on the state of the world.

Taken together, these results suggest that any and all subjective beliefs are affected by wishful thinking bias, and that the bias may well be sufficiently strong to

materially affect economically important decisions. High stakes decisions in financial markets are a case in point, as they involve probability assessments in situations characterized by high stakes and high subjective uncertainty—both of which are conducive to the presence of an economically significant bias.

In interpreting this conclusion, it is important to bear in mind that decision makers in high stakes situations have an incentive to invest in quality information in order to reduce the uncertainty in their beliefs. Since the strength of the bias depends on the degree of subjective uncertainty, quality information will not only reduce the variance in beliefs, but would also (perhaps unintentionally) reduce the magnitude of the bias. The degree to which wishful thinking is likely to affect high stakes decisions is therefore dependent on decision makers' ability to reduce uncertainty before making their choices.

One way to assess the degree of uncertainty is to examine the beliefs of informed experts. In many important decision making environments (financial markets, corporate decision making, politics, war) informed experts commonly disagree. The failure of experts to come to anything approaching consensus suggests the existence of a substantial level of irreducible uncertainty. When that is the case, there is evidently significant potential for wishful thinking to materially affect decisions.

The present paper describes one particular experiment on one particular group of subjects. While the main conclusions are strongly statistically significant, it would clearly be important to see whether the results can be replicated by other researchers and in other decision making environments. Another important limitation in interpreting the results of the experiment is the limited range of theories under consideration. While I am not aware of any other non ad hoc theory that can explain the results of the experiment, it is important to emphasize that if such a theory were to be offered, it may significantly change the interpretation of the experiment's results.

References

- Akerlof, G. and Dickens, W. (1982). The Economic Consequences of Cognitive Dissonance, *American Economic Review* **72**(3): 307–319.
- Armantier, O. and Treich, N. (2010). Eliciting beliefs: Proper scoring rules, incentives, stakes and hedging, *TSE Working Papers*, Toulouse School of Economics.

- Babcock, L. and Loewenstein, G. (1997). Explaining Bargaining Impasse: The Role of Self-Serving Biases, *The Journal of Economic Perspectives* **11**(1): 109–126.
- Babcock, L., Loewenstein, G., Issacharoff, S. and Camerer, C. (1995). Biased Judgments of Fairness in Bargaining, *American Economic Review* **85**: 1337–1343.
- Barber, B. and Odean, T. (2001). Boys Will be Boys: Gender, Overconfidence, and Common Stock Investment, *Quarterly Journal of Economics* **116**(1): 261–292.
- Benabou, R. and Tirole, J. (2002). Self-confidence and personal motivation, *Quarterly Journal of Economics* **117**(3): 871–915.
- Benoît, J. and Dubra, J. (2011). Apparent Overconfidence, *Econometrica* **79**(5): 1591–1625.
- Blanco, M., Engelmann, D., Koch, A. and Normann, H. (2008). Belief elicitation in experiments: is there a hedging problem?, *Experimental Economics* pp. 1–27.
- Brunnermeier, M. and Parker, J. (2005). Optimal Expectations, *American Economic Review* **95**(4): 1092–1118.
- Camerer, C. and Lovallo, D. (1999). Overconfidence and excess entry: An experimental approach, *American Economic Review* **89**(1): 306–318.
- Caplin, A. and Leahy, J. (2001). Psychological Expected Utility Theory and Anticipatory Feelings, *Quarterly Journal of Economics* **116**(1): 55–80.
- Carrillo, J. and Mariotti, T. (2000). Strategic Ignorance as a Self-Disciplining Device, *Review of Economic Studies* **67**(3): 529–544.
- Compte, O. and Postlewaite, A. (2004). Confidence-Enhanced Performance, *American Economic Review* **94**(5): 1536–1557.
- Cover, T. and Thomas, J. (1991). *Elements of information theory*, Vol. 1, Wiley Online Library.
- Edwards, R. and Magee, J. (2010). *Technical analysis of stock trends*, Snowball Publishing.
- Eil, D. and Rao, J. (2011). The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself, *American Economic Journal: Microeconomics* **3**(2): 114–138.
- Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments, *Experimental Economics* **10**(2): 171–178.

- Gneiting, T. and Raftery, A. (2007). Strictly proper scoring rules, prediction, and estimation, *Journal of the American Statistical Association* **102**(477): 359–378.
- Good, I. (1952). Rational decisions, *Journal of the Royal Statistical Society. Series B (Methodological)* pp. 107–114.
- Johnson, D. H. and Sinanovic, S. (2001). Symmetrizing the kullback-leibler distance, *IEEE Transactions on Information Theory* **1**(1): 1–10.
- Kőszegi, B. (2006). Ego Utility, Overconfidence, and Task Choice, *Journal of the European Economic Association* **4**(4): 673–707.
- Knight, F. (1921). *Risk, Uncertainty and Profit*, Houghton Mifflin.
- Knox, R. and Inkster, J. (1968). Postdecision dissonance at post time, *Journal of Personality and Social Psychology* **8**(4): 319–23.
- Kunda, Z. (1990). The case for motivated reasoning, *Psychological Bulletin* **108**(3): 480–498.
- Loewenstein, G., Issacharoff, S., Camerer, C. and Babcock, L. (1993). Self-serving assessments of fairness and pretrial bargaining, *The Journal of Legal Studies* **22**(1): 135–159.
- Lundeberg, M., Fox, P., Brown, A. and Elbedour, S. (2000). Cultural influences on confidence: Country and gender, *Journal of Educational Psychology* **92**(1): 152.
- Malmendier, U. and Tate, G. (2005). CEO overconfidence and corporate investment, *Journal of Finance* **60**(6): 2661–2700.
- Malmendier, U. and Tate, G. (2008). Who makes acquisitions? CEO overconfidence and the market’s reaction, *Journal of Financial Economics* **89**(1): 20–43.
- Mayraz, G. (2011). Priors and Desires—A Model of Payoff-Dependent Beliefs, *CEP Discussion Paper 1047*, Centre for Economic Performance, London School of Economics.
- Mijović-Prelec, D. and Prelec, D. (2010). Self-deception as self-signalling: a model and experimental evidence, *Philosophical Transactions of the Royal Society B: Biological Sciences* **365**(1538): 227.
- Miller, D. and Ross, M. (1975). Self-serving biases in the attribution of causality: Fact or fiction, *Psychological Bulletin* **82**(2): 213–225.

- Mullainathan, S. and Washington, E. (2009). Sticking With your Vote: Cognitive Dissonance and Political Attitudes, *American Economic Journal: Applied Economics* **1**(1): 86–111.
- Murphy, J. J. (1999). *Technical analysis of the financial markets*, New York Institute of Finance.
- Nisbett, R. and Ross, L. (1980). *Human inference: Strategies and shortcomings of social judgment*, Prentice-Hall Englewood Cliffs, NJ.
- Olsen, R. (1997). Desirability bias among professional investment managers: Some evidence from experts, *Journal of Behavioral Decision Making* **10**(1): 65–72.
- Park, Y. and Santos-Pinto, L. (2010). Overconfidence in tournaments: evidence from the field, *Theory and Decision* **69**(1): 143–166.
- Santos-Pinto, L. and Sobel, J. (2005). A model of positive self-image in subjective assessments, *American Economic Review* pp. 1386–1402.
- Stone, R. (2006). Self-handicapping, Tournaments and Ego Utility, Mimeo.
- Van den Steen, E. (2004). Rational Overoptimism (and Other Biases), *American Economic Review* **94**(4): 1141–1151.
- Yariv, L. (2002). I'll See it when I Believe it: A Simple Model of Cognitive Consistency, Mimeo.

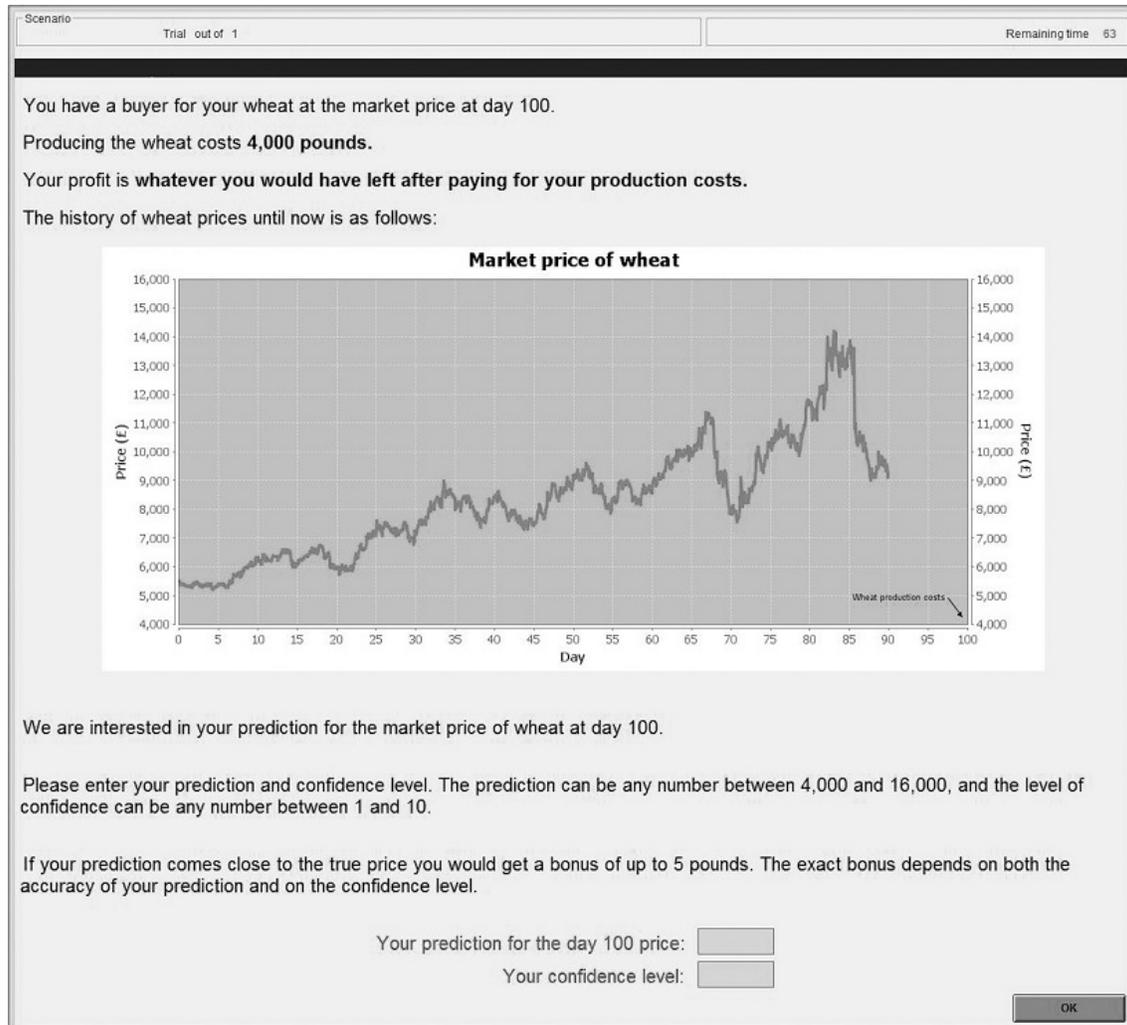


Figure 1: The interface of the *Farmers* treatment with a maximum accuracy bonus of £5. The interface of the *Bakers* treatment was similar, except: (a) the first three lines were: “*You have a buyer for £16,000 worth of bread from your bakery. At day 100 you will get the money from the order, and will have to use some of it to buy wheat at the market. Your profit is whatever you would have left after paying for the wheat.*”, and (b) instead of an arrow on the chart pointing to £4,000 with the label “*Wheat production costs*”, there was an arrow pointing to £16,000 with the label “*The price you would get for your bread*”.

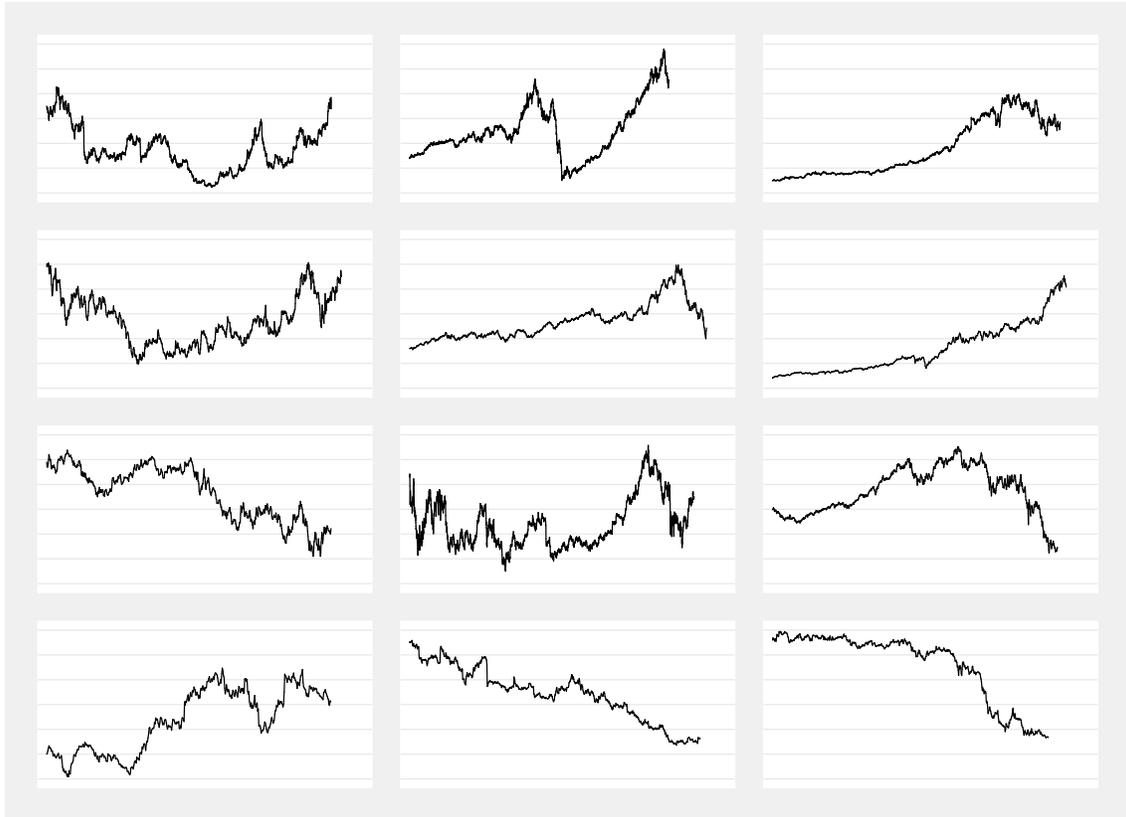


Figure 2: The charts used in the 12 earning periods. The x -axis represents time, ranging from day 0 to day 100, and the y -axis represents price, ranging from £4,000 to £16,000. The data for the charts were adapted from historical equity price data, shifted and scaled to fit into a uniform range. Figure 1 shows how these charts were presented to subjects.

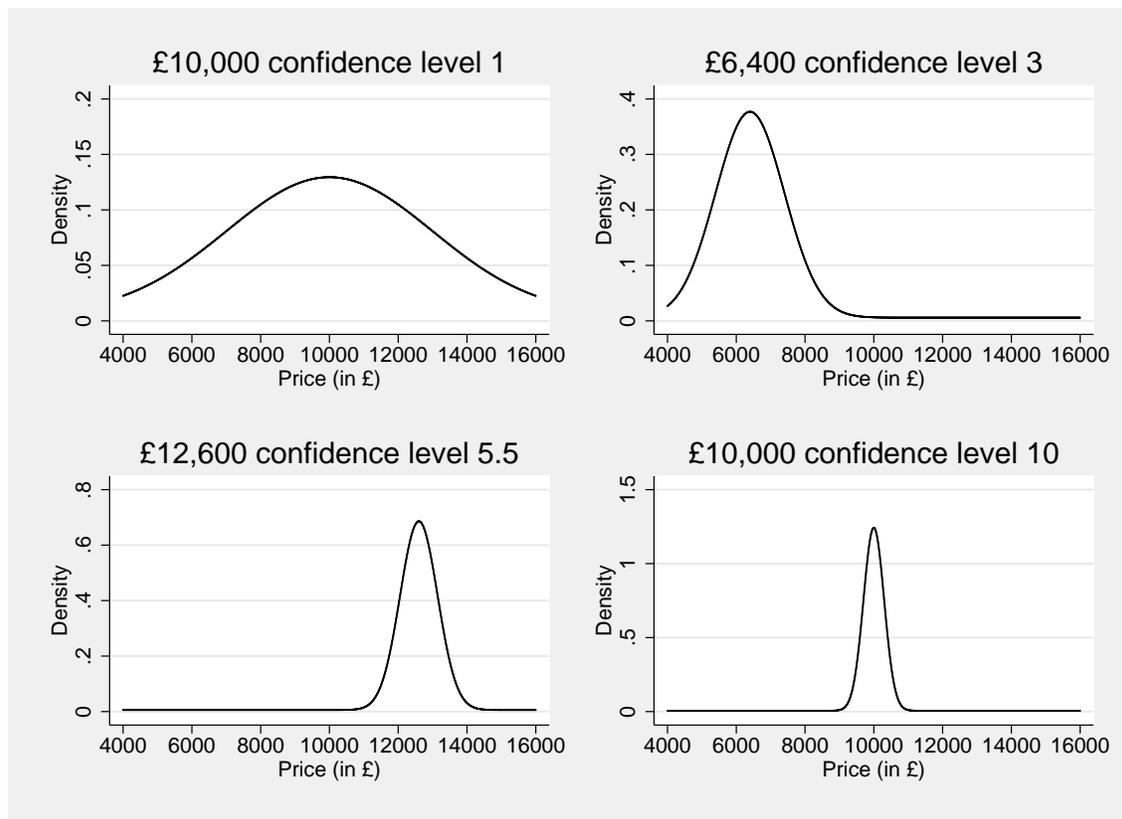


Figure 3: The examples of distributions used in the instructions. Each distribution is characterized by a prediction and a confidence level. These examples were used in explaining the prediction elicitation procedure. They were particularly useful in establishing a reference for the 1-10 scale that was used in reporting the subject's confidence in her prediction.

Table 3: Wishful thinking bias and comparative statics. The table reports the estimated bias in different sub-samples and statistical tests of related hypotheses.

	Sample	Estimated bias ^a	Observations ^b
	All subjects	452*** (s.e. 123)	1584 (132)
	<i>negative ?</i>	$p < 0.0002$	
Cost of bias	Accuracy bonus: <i>low</i> (£1)	298** (s.e. 164)	816 (68)
	Accuracy bonus: <i>medium</i> (£2)	569** (s.e. 328)	300 (25)
	Accuracy bonus: <i>high</i> (£5)	645*** (s.e. 210)	468 (39)
	<i>low = medium = high ?</i>	$p < 0.4026$	
	<i>low = 2 · medium = 5 · high ?</i>	$p < 0.0140^c$	
Degree of subjective uncertainty	Chart uncertainty: <i>low</i>	269** (s.e. 127)	792 (66)
	Chart uncertainty: <i>high</i>	635*** (s.e. 166)	792 (66)
	<i>low > high ?</i>	$p < 0.0142$	
	Within chart variance: <i>low</i>	227** (s.e. 113)	792 (66)
	Within chart variance: <i>high</i>	677*** (s.e. 175)	792 (66)
	<i>low > high ?</i>	$p < 0.0034$	
Stakes in the value of the day 100 price	Stakes: <i>low</i> (50p)	260 (s.e. 289)	288 (24)
	Stakes: <i>standard</i> (£1)	495*** (s.e. 135)	1296 (108)
	<i>standard ≤ 2 · low ?</i>	$p < 0.2313^d$	
	<i>standard = 2 · low ?</i>	$p < 0.9668$	
Confidence in ability to predict prices	Average confidence: <i>low</i>	276* (s.e. 174)	792 (66)
	Average confidence: <i>high</i>	628*** (s.e. 169)	792 (66)
	<i>low > high ?</i>	$p < 0.0732$	
	Prices predictable? <i>no</i>	292** (s.e. 174)	792 (66)
	Prices predictable? <i>yes</i>	613*** (s.e. 174)	792 (66)
	<i>no > yes ?</i>	$p < 0.0997$	
Demographics	Males	411** (s.e. 187)	600 (50)
	Females	477*** (s.e. 166)	984 (82)
	<i>same ?</i>	$p < 0.7956$	

^a Robust standard errors in parentheses. Statistical significance indicators: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

^b An individual observation refers to the prediction of a given subject in a given chart. Clustering is by subjects. The number of clusters is in parentheses.

^c If the regression is restricted to the sessions with standard stakes the test p -values are 0.5094 and 0.0171 respectively.

^d If the regression is restricted to the sessions with a low maximum bonus the test p -values are 0.7620 and 0.4269 respectively.

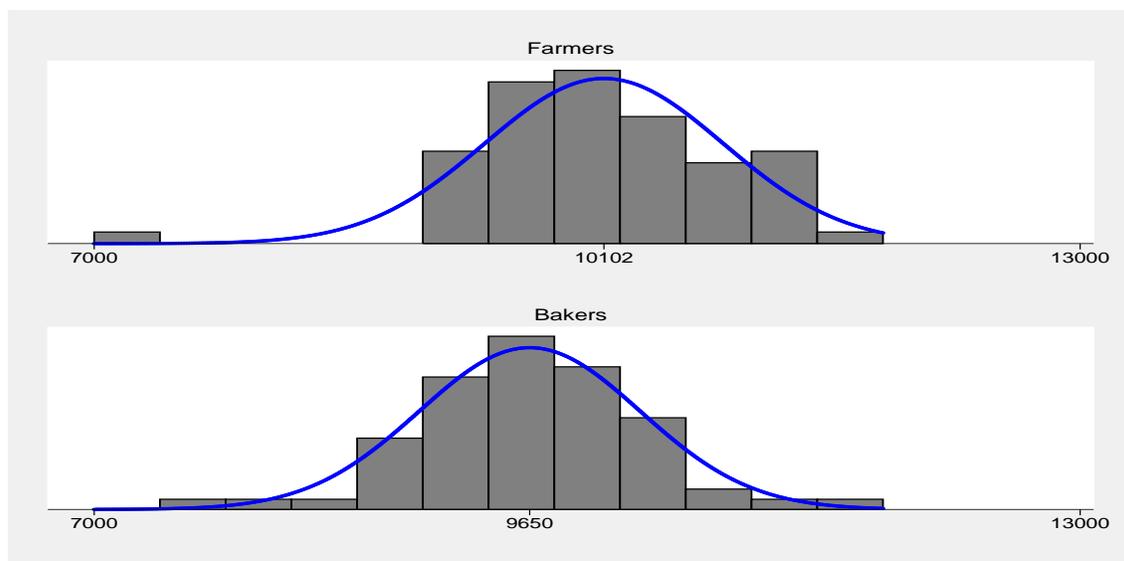


Figure 4: Histogram of the mean predictions made by *Farmers* and *Bakers*. A normal distribution curve was fitted to both histograms. The mean prediction was 10102 and 9650 respectively. 16 of the 20 subjects making the highest (lowest) mean predictions were *Farmers* (*Bakers*).

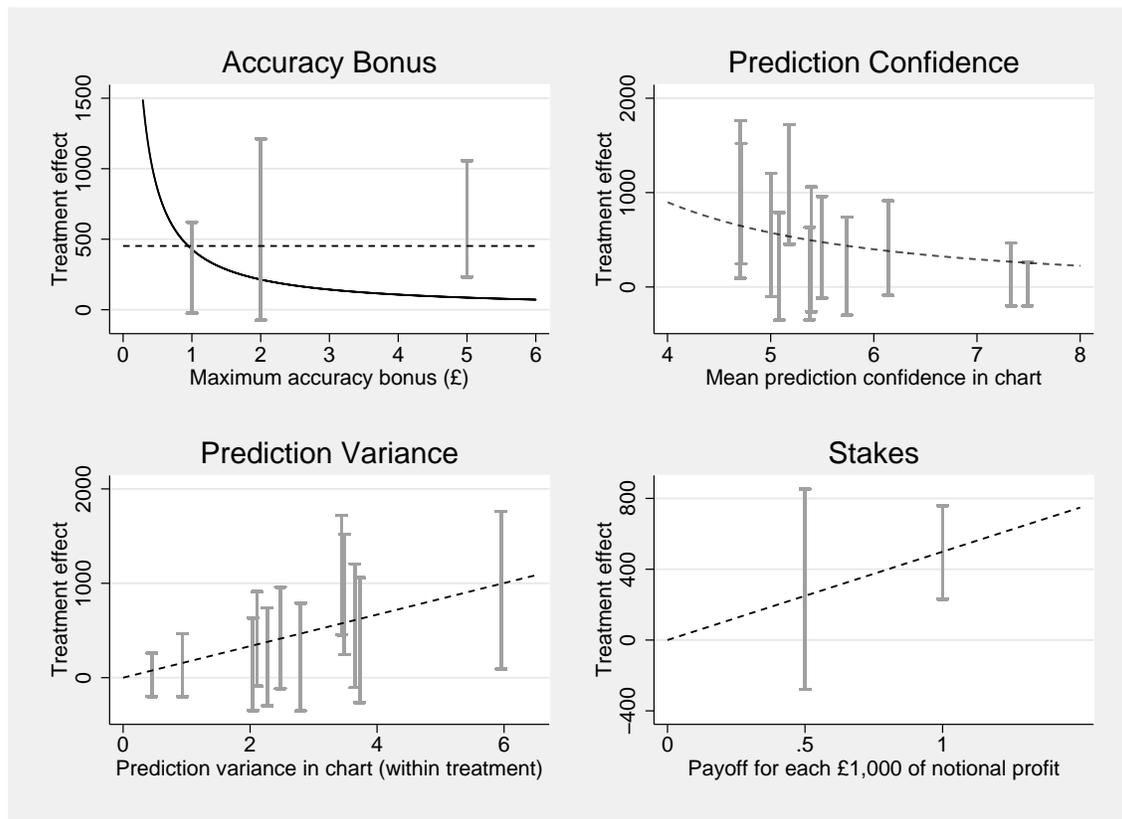


Figure 5: The comparative statics of wishful thinking bias. The panels show a 95 percent confidence interval for difference in predictions between *Farmers* and *Bakers* (the treatment effect) in different subsamples. The first panel shows the comparative statics of the cost of holding wrong beliefs, represented by the maximum accuracy bonus. The solid hyperbolic line represents the best fit for the Optimal Expectations model, and the dashed horizontal line that of Priors and Desires. The second panel shows the bias in a chart against the mean confidence in predictions for that chart. The curve is fitted to the inverse of the square of the mean confidence level. The third panel shows the bias in a chart against the mean within group predictions variance. The dashed line is a linear fit through the origin. Finally, the fourth panel shows the comparative statics of the stakes, the x -axis representing the amount in pounds that a subject receives for each £1,000 of notional profit. The dashed line is a linear fit through the origin.

CENTRE FOR ECONOMIC PERFORMANCE
Recent Discussion Papers

1091	Francesco Caselli Andrea Tesei	Resource Windfalls, Political Regimes, and Political Stability
1090	Keyu Jin Nan Li	Factor Proportions and International Business Cycles
1089	Yu-Hsiang Lei Guy Michaels	Do Giant Oilfield Discoveries Fuel Internal Armed Conflicts?
1088	Brian Bell John Van Reenen	Firm Performance and Wages: Evidence from Across the Corporate Hierarchy
1087	Amparo Castelló-Climont Ana Hidalgo-Cabrillana	The Role of Educational Quality and Quantity in the Process of Economic Development
1086	Amparo Castelló-Climont Abhiroop Mukhopadhyay	Mass Education or a Minority Well Educated Elite in the Process of Development: the Case of India
1085	Holger Breinlich	Heterogeneous Firm-Level Responses to Trade Liberalisation: A Test Using Stock Price Reactions
1084	Andrew B. Bernard J. Bradford Jensen Stephen J. Redding Peter K. Schott	The Empirics of Firm Heterogeneity and International Trade
1083	Elisa Faraglia Albert Marcet Andrew Scott	In Search of a Theory of Debt Management
1082	Holger Breinlich Alejandro Cuñat	A Many-Country Model of Industrialization
1081	Francesca Cornaglia Naomi E. Feldman	Productivity, Wages and Marriage: The Case of Major League Baseball
1080	Nicholas Oulton	The Wealth and Poverty of Nations: True PPPs for 141 Countries
1079	Gary S. Becker Yona Rubinstein	Fear and the Response to Terrorism: An Economic Analysis
1078	Camille Landais Pascal Michaillat Emmanuel Saez	Optimal Unemployment Insurance over the Business Cycle
1077	Klaus Adam Albert Marcet Juan Pablo Nicolini	Stock Market Volatility and Learning

1076	Zsófia L. Bárány	The Minimum Wage and Inequality - The Effects of Education and Technology
1075	Joanne Lindley Stephen Machin	Rising Wage Inequality and Postgraduate Education
1074	Stephen Hansen Michael McMahon	First Impressions Matter: Signalling as a Source of Policy Dynamics
1073	Ferdinand Rauch	Advertising Expenditure and Consumer Prices
1072	Alberto Galasso Mark Schankerman Carlos J. Serrano	Trading and Enforcing Patent Rights
1071	Daniel Hale John Coleman Richard Layard	A Model for the Delivery of Evidence-Based PSHE (Personal Wellbeing) in Secondary Schools
1070	David Marsden	Individual Voice in Employment Relationships: A Comparison Under Different Forms of Workplace Representation
1069	Ross Levine Alexey Levkov Yona Rubinstein	Racial Discrimination and Competition
1068	Klaus Adam Albert Marcet	Internal Rationality, Imperfect Market Knowledge and Asset Prices
1067	Yanhui Wu	A Simple Theory of Managerial Talent, Pay Contracts and Wage Distribution
1066	Yanhui Wu	Managerial Incentives and Compensation in a Global Market
1065	Nicholas Bloom Helena Schweiger John Van Reenen	The Land that Lean Manufacturing Forgot? Management Practices in Transition Countries
1064	Klaus Adam Pei Kuang Albert Marcet	House Price Booms and the Current Account
1063	Stephen Hansen Michael McMahon	How Experts Decide: Identifying Preferences versus Signals from Policy Decisions
1062	Paul Dolan Daniel Fujiwara Robert Metcalfe	A Step Towards Valuing Utility the Marginal and Cardinal Way