

Wishful Thinking

Guy Mayraz*

July 23, 2013

Abstract

A Bayesian model of wishful thinking is proposed. Agents start out with a state-dependent endowment, and form beliefs as if Nature was more likely to have chosen a state in which their payoff is high. In an associated experiment, subjects were given opposite stakes in the future price of a financial asset: some stood to gain from high prices, and others from low prices. Both groups were also paid for the accuracy of their predictions. Despite incentives for hedging, subjects gaining from high prices predicted higher prices than subjects gaining from low prices. Increasing the accuracy bonus did not result in a smaller bias. The proposed model is consistent with these results and with existing field evidence suggesting that wishful thinking is a pervasive bias that can affect decisions large and small.

JEL classification: D01,D03,D80,D81,D83,D84,G11.

Keywords: wishful thinking, optimism, pessimism, cognitive dissonance, reference-dependent beliefs, reference-dependent preferences.

*Department of Economics, University of Melbourne. guy.mayraz@unimelb.edu.au. 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, and to seminar participants at the Hebrew University, London School of Economics, MIT, Oxford, Princeton, Tel-Aviv, UBC, UC Berkeley, and UC Santa Barbara. I gratefully acknowledge financial support from the Russell Sage Foundation Small Grants Program in Behavioral Economics and the John Fell OUP Research Fund.

1 Introduction

People exhibit wishful thinking if they are more likely to believe something is true simply because they *want* it to be true. If a wishful-thinker wants to have high ability, to avoid bad outcomes, to win conflicts, and to always do the right thing, that person would be overconfident, underestimate the risks she is exposed to, expect to beat her opponents, and believe that justice is with her.

Wishful thinking is, in principle, a powerful force with a substantial impact on important decisions. But how important is it in practice? The belief patterns predicted by wishful thinking have been observed in many areas of economics,¹ consistent with the idea that wishful thinking is a pervasive bias that affects decisions large and small. However, while this evidence is certainly suggestive, it is nearly always possible to find alternative explanations for each particular finding.²

Wishful thinking is presumably irrelevant in situations probabilities are given, or are readily calculated. For example, it stretches credulity that a person would really believe a coin would land on Heads simply because she stands to win a large sum of money if it does. We would not be surprised if a person in this situation has a feeling that the coin would land on Heads, but since she knows the correct probabilities, we would expect this knowledge to override her intuition. However, as [Knight \(1921\)](#) noted nearly a century ago, most of the situations economists care are rather different. Beliefs about such things as the success of a business venture or the future course of house

¹[Babcock and Loewenstein \(1997\)](#) link the low frequency of pretrial bargains to a tendency by both parties to believe that they would win if the case ends up in court. [Olsen \(1997\)](#) finds evidence for optimistically biased beliefs among professional investment managers. [Camerer and Lovallo \(1999\)](#) link excess entry into competitive markets to overconfidence over relative ability. [Malmendier and Tate \(2008\)](#) argue that managerial overconfidence is responsible for corporate investment distortions. [Cowgill et al. \(2009\)](#) find optimistic bias in corporate prediction markets. [Park and Santos-Pinto \(2010\)](#) provide field evidence for overconfidence in tournaments. [Hoffman \(2011a,b\)](#) finds that truck drivers are optimistically biased about their productivity (and hence their pay), resulting in an inefficient failure to switch jobs.

²For example, in [Babcock and Loewenstein \(1997\)](#) subjects in the role of plaintiff came to expect higher penalties than subjects in the role of defendant, even though both groups of subjects were exposed to the same case materials. However, subjects had to argue their side with the other party, which may have caused them to focus their reading on arguments favoring their case. Their beliefs could thus have arisen from a failure to correct for this selective attention, rather than from a general wishful thinking bias.

prices are not constrained by given or readily calculable probabilities, and decision makers affected by wishful thinking may well remain oblivious to its effects. Needless to say, if a person cannot tell that her beliefs are biased, she cannot prevent the bias from affecting her decisions.

But why should we expect wishful thinking in the first place? One important idea is that people may wish to deceive themselves. In the presence of other biases, wishful thinking can sometimes help a person obtain desirable outcomes. Wishful-thinking over the likelihood of success (over-confidence) can motivate a person who would otherwise underinvest effort because of weak willpower (Benabou and Tirole, 2002), make it easier for people to convince others of their ability (ibid),³ and protect them from negative emotions that may be detrimental to performance (Compte and Postlewaite, 2004). Alternatively, wishful thinking simply *feels* good, and may therefore be desirable in itself (Akerlof and Dickens, 1982; Carrillo and Mariotti, 2000; Caplin and Leahy, 2001; Benabou and Tirole, 2002; Brunnermeier and Parker, 2005).

In addition to a suitable motive, self-deception also requires a suitable technology. In models of instrumental self-deception (Benabou and Tirole, 2002; Compte and Postlewaite, 2004) the technology is selective memory: repressing the memory of failures and ensuring that successes are well remembered.⁴ Models of hedonic self-deception (Akerlof and Dickens, 1982; Brunnermeier and Parker, 2005) simply assume that people can choose the beliefs of their future selves at will.⁵

Instrumental self-deception counteracts the potential impact of other biases. Consequently, even when the effect on beliefs is substantial, the net effect on *choices* (over a model with no behavioral biases) can be small, or even disappear altogether.⁶ Models of hedonic self-deception assume no other behavioral biases, and can potentially result in a substantial impact on choices. However, these models assume that choices in the absence of wishful-thinking are optimal. Wishful thinking is therefore pulling choices

³The behavioral bias in this case is a difficulty to lie about one's private information.

⁴Effectiveness depends on whether the tampered memory is interpreted naively or sophisticatedly (Benabou and Tirole, 2002).

⁵There is also a class of models with hedonic preferences, which assume no self deception technology apart from control over the acquisition of signals (Carrillo and Mariotti, 2000; Kőszegi, 2006). The interest in such models comes from the ability to manipulate the higher moments of the distribution of beliefs (the mean is fixed by the law of iterated expectations).

⁶It would only be really zero if there exist beliefs that fully counteract the effect of other behavioral biases, and the technology for deceiving oneself is perfect.

away from optimality, and we should only expect to see substantial wishful thinking if its impact on choices is small. The wishful thinking that results from self-deception thus has only limited obvious implications for choices over a standard model of decision making with no behavioral biases. There are therefore good reasons to consider a different modeling approach if we are interested in the possibility that wishful thinking has a significant impact on high stakes decisions.

One alternative is that wishful thinking arises earlier in the decision making process, before people even think about the choices that lie ahead of them. Perhaps judgments of subjective desirability (do I *want* this to be true?) leak into judgments of subjective likelihood (is this *likely* to be true?), so that desirable events are perceived as more likely. Wishful thinking isn't then a choice people make, but is rather a fundamental feature of how they perceive the world. An important implication is that there is then no direct relationship between the magnitude of the bias in a person's beliefs and the importance of the decisions that depend on them.⁷ Decisions large and small can be equally affected by wishful thinking.

One contribution of the present paper is a model, *Priors and Desires*, which formalizes this idea. The perceived likelihood of a state ω , in which the decision maker has a stake $r(\omega)$ (in utility terms) is given by the following equation, where p_0 represents the beliefs the decision maker would have held if she were indifferent between all states, and ψ is a parameter that characterizes the decision maker:

$$p_r(\omega) \propto p_0(\omega)e^{\psi r(\omega)} \tag{1}$$

A decision maker whose beliefs are represented by Equation 1 is a wishful thinker if $\psi > 0$. Other things being equal, the higher the payoff in a state, the higher its subjective probability. Sharper predictions are obtained if we assume that p_0 coincides with rational expectations, but the model does not require this assumption.

Equation 1 is formally equivalent to Bayes rule with $\psi r(\omega)$ as the log likelihood in state ω . Decision makers believe that Nature chose the state of the world with their interests in mind. Given this belief, they observe what they have to gain or lose if an event obtains, and use it as evidence in judging its likelihood.

Since the payoff in a state functions as information, the impact on beliefs

⁷Holding information constant, there is no relationship between the magnitude of the bias and the importance of subsequent decisions.

depends on the degree of uncertainty and on what the decision maker has at stake, but not on the importance of the decisions that depend on getting beliefs right. It follows that wishful thinking is a pervasive phenomenon that exists whenever a decision maker approaches a choice situation with some existing stake in what is or isn't the case, and which does not diminish simply because of their detrimental effect on choices.

In addition to wishful thinking or *optimism*, the *Priors and Desires* model can also be used to capture *pessimism* (if $\psi < 0$). Moreover, since beliefs in the model depend on what the decision maker has at stake, a change in these stakes causes a change in beliefs, even in the absence of any change in relevant information. This can be used to model *cognitive dissonance*.

The second contribution of the paper is a controlled experiment designed to provide an unambiguous test of wishful thinking, to study its comparative statics, and to enable us to say something about the mechanism behind it. Student subjects were randomly assigned into one of two groups: *Farmers*, whose payoff was increasing in the price of wheat, and *Bakers*, whose payoff was decreasing in this price. Subjects in both groups were shown charts of historical wheat prices, and were given a performance bonus for making good predictions of what the future price would be. Both groups therefore had the same incentive to reveal their beliefs over future wheat prices, but opposite stakes in what this price would be. Consistent with wishful thinking, and despite the possible effect of hedging, *Farmers* made significantly higher predictions than *Bakers*, and the null hypothesis of no bias was strongly rejected.⁸

This setting provides no instrumental benefit to wishful thinking, and we can therefore rule out instrumental self-deception as a possible explanation. Moreover, since all relevant information is available when predictions are made, there is no opportunity for memory manipulation. In the absence of a suitable technology for self-deception, hedonic self deception is also an unlikely mechanism. As a further test of the underlying mechanism, the size of the accuracy bonus was altered between sessions. The bigger the bonus, the bigger the expected cost of any given level of bias. Hence, if the bias is caused by hedonic self-deception, we should expect it to decrease with the size of the accuracy bonus. However, no decrease in the magnitude of the

⁸*p*-value of either 0.0016 or 0.0001, depending on how outliers are treated. Risk aversion over small stakes could have led subjects to intentionally bias their reports, but such hedging incentives would pull predictions in the opposite direction to wishful thinking.

bias was observed, and the specific prediction of the best known hedonic self-deception model, *Optimal Expectations* (Brunnermeier and Parker, 2005), was formally rejected.⁹ While self-deception has difficulty accounting for the bias in this situation, the findings are consistent with a model such as *Priors and Desires*, which predicts wishful thinking whenever decision makers have a stake in what is or isn't the case.

Further evidence is provided by the comparative statics of confidence. The more confident a person is in her ability to predict the day 100 prices, the more she has to lose from making biased predictions. A subject's confidence was defined by the average confidence rating in her predictions across all charts, and a second confidence measure was obtained by asking subjects in the post experiment questionnaire whether they believe that prices in financial markets are generally predictable. In both cases the opposite result was obtained: a general belief in the predictability of prices and confidence in predictions were *positively* correlated with the degree of bias. This result is explained by noting that subjects have a stake in their ability to make good predictions. Highly optimistic subjects (having a high value of ψ in Equation 1) are both more biased about their ability to predict prices *and* more biased in each individual prediction.

The experiment also tested how the magnitude of the bias varies with the degree of subjective uncertainty across different prediction situations. Both *Priors and Desires* and *Optimal Expectations* predict a positive relationship (the intuition is that greater subjective uncertainty leaves more scope for wishful thinking to affect beliefs). This prediction was tested by assigning charts into two groups according to the average confidence rating across all subjects (a chart involves a great deal of subjective uncertainty if subjects lack confidence when making their predictions). As predicted, the magnitude of the bias was significantly greater in high uncertainty (low confidence) charts.¹⁰

The remainder of the paper is organized as follows. Section 2 introduces the *Priors and Desires* model. Section 3 describes the experiment, and develops the predictions of *Priors and Desires* and of *Optimal Expectations*. Section 4 presents the results of the experiment, and Section 5 concludes.

⁹ p -value of approximately 0.01. See Section 4 for details. The test assumes risk neutrality over small stakes.

¹⁰Each particular confidence rating is a joint function of the subject making the prediction and the prediction situation (the identity of the chart). The greatest bias is for confident subjects in low confidence (high uncertainty) charts.

2 The *Priors and Desires* model

Priors and Desires is a model of choice under uncertainty that allows for the possibility that a person’s subjective beliefs (her ‘priors’) are affected by what she has at stake (her ‘desires’). An optimist (or wishful thinker) is more likely to believe something is true simply because she has a stake in it being true; a pessimist is biased in the opposite direction: the more she wants something to be true, the *lower* its subjective probability. Choice is otherwise standard: decision makers take their beliefs at face value, and maximize subjective expected utility in their decisions.

The setup is simple. At $t = 0$ Nature chooses some particular state ω^* . Nature then reveals to the decision maker (i) a *signal* about ω^* and (ii) the decision maker’s *initial stakes* in what ω^* is. The decision maker then has a choice to make, adding to her initial stakes. Finally, ω^* is revealed, and the decision maker obtains the payoff that her initial stakes and choice yield in ω^* .

In the experiment, states correspond to possible values for the final wheat price. *Farmers (Bakers)* start out with a stake in high (low) wheat prices, as well as a signal in the form of a chart of historical prices. They then choose what price to bet on. Finally, the true price is revealed, and subjects receive their combined payoff, which depends on the true price of wheat.

Formally, let Ω denote a finite¹¹ set of states, and let S denote a set of (subjective) signals about Ω . A *payoff-function* $f : \Omega \rightarrow \mathbb{R}$ is a mapping assigning to each state $\omega \in \Omega$ a real number $f(\omega)$, representing the payoff in utility terms that is obtained if ω is realized. Let $F = \{f : \Omega \rightarrow \mathbb{R}\}$ denote the set of all payoff-functions. Payoff-functions are used to represent both the initial (or *reference*) stakes r and the choice c . Utility is assumed to be additive, so that the combined payoff in state ω is $r(\omega) + c(\omega)$. Timing is as follows: at $t = 1$ the decision maker observes the initial stakes $r \in F$ and a signal $s \in S$. At $t = 2$ she chooses some alternative c from a choice set $C \subseteq F$. Finally, at $t = 3$ some particular state ω^* is realized, and the decision maker obtains the payoff $r(\omega^*) + c(\omega^*)$.

A standard decision maker with a subjective prior probability measure q over $\Omega \times S$ would observe the signal s , and choose c to maximize expected payoff according to $q(\cdot|s)$. The reference stakes r are irrelevant to her choice. A *Priors and Desires* decision maker is different. Such a decision maker is

¹¹This restriction is purely for expositional purposes. It is relaxed in Section 2.1.

characterized by the combination of a probability measure p_0 and a parameter $\psi \in \mathbb{R}$, called the *coefficient of relative optimism*. She chooses c to maximize expected payoff not according to $p_0(\cdot|s)$, but according to a probability measure $p_r(\cdot|s)$ defined by the following equation:

$$p_r(\omega|s) \propto p_0(\omega|s)e^{\psi r(\omega)}. \quad (2)$$

To understand Equation 2 consider first the special case where the decision maker starts out with nothing at stake, so that for any two states ω and ω' , $r(\omega) = r(\omega')$. The term $e^{\psi r(\omega)}$ is then independent of ω , and can be dropped out of the equation, with the result that $p_r(\cdot|s) = p_0(\cdot|s)$. More generally, r does vary between states, and $p_r(\cdot|s) \neq p_0(\cdot|s)$. Other things being equal, if ψ is positive (negative) $p_r(\cdot|s)$ is higher in states in which r is higher (lower). A positive value of ψ therefore represents *optimistic bias* (or wishful thinking), and a negative value represents *pessimistic bias*. Finally, note that if $\psi = 0$ the reference stakes are irrelevant. Such decision makers are *realists*, and for them $p_r(\cdot|s) = p_0(\cdot|s)$ for all r and s . A standard subjective expected utility maximizing decision maker is equivalent to a *Priors and Desires* decision maker with a coefficient of relative optimism of zero.

When comparing the likelihood of one state ω relative to another state ω' , we can take logs to obtain the following expression:

$$\log \frac{p_r(\omega|s)}{p_r(\omega'|s)} = \log \frac{p_0(\omega|s)}{p_0(\omega'|s)} + \psi [r(\omega) - r(\omega')]. \quad (3)$$

The subjective log odds-ratio of a *Priors and Desires* agent equals what it would have been had she been indifferent between the two states ($r(\omega) = r(\omega')$) plus a term which depends linearly on the payoff difference between them. If the agent is an optimist ($\psi > 0$) the relative likelihood of the more desirable (higher utility) state is shifted upwards, and if the agent is a pessimist ($\psi < 0$) the opposite is true. The higher ψ is in absolute terms, the greater the bias. In the limit of $\psi \rightarrow \pm\infty$ the agent is certain that the most desirable (least desirable) state is true.

An important feature of the *Priors and Desires* model is that its equations are formally identical to Bayes Rule, with $\psi r(\omega)$ as the log likelihood in state ω . Consider the extended state space $\Omega \times S \times F$, and define a probability measure q as follows for all $\omega \in \Omega$, $s \in S$ and $r \in F$:

$$q(\omega, s, r) = p_0(\omega)p_0(s|\omega)e^{\psi r(\omega)}. \quad (4)$$

Given this definition, the *Priors and Desires* probability measure p_r coincides with the outcome of Bayesian updating on q :

$$p_r(\cdot|s) = q(\cdot|r, s). \tag{5}$$

It is thus possible to interpret the equations of the model as Bayesian updating by decision makers who believe Nature had their interests in mind when choosing ω^* . An optimistic decision maker believes that Nature is benevolent, and that states in which she obtains a relatively high utility are more likely to have been chosen by Nature. A pessimistic decision maker makes the opposite inference. Decision makers observe their initial stakes, and *infer* what state Nature was likely to have selected. The initial stakes function as a second signal.

It is not necessary to assume that decision makers consciously believe that their interests are informative about the state of the world. An alternative interpretation is that this belief is unconscious, but that it nonetheless affects the subjective probability judgments that decision makers arrive at. Decision maker are unaware of this property of their subjective judgment, and proceed to use its biased output in their decisions.¹²

The *Priors and Desires* model also offers a novel way of thinking about cognitive dissonance. The essence of cognitive dissonance is that new information that alters people’s interests shifts their beliefs in the direction of the new interests, even when there is no new evidence that could conventionally explain the change in beliefs. For example, in a classical cognitive dissonance experiment [Festinger and Carlsmith \(1959\)](#) found that students who were asked to recommend a boring task to another student later rated the task as much less boring than did other students. Their explanation is that recommending the task causes dissonance with the belief that the task is boring, and that the dissonance is resolved by altering this belief, so that it fits better with the action. Alternatively, we can view the act of recommending the task as giving the students a stake in the task being relatively interesting. The students take this stake as new information about Nature’s choice of state, and update their beliefs accordingly.

This model has some resemblance to the game against Nature interpretation of ambiguity aversion ([Maccheroni et al., 2006](#)). However, in the ambiguity aversion case, Nature moves *after* the decision maker, and has the

¹²These two interpretations have different implications for the decision maker’s beliefs over her future choices, but in the one period setting they are observationally equivalent. It is similarly possible to interpret self-deception as either conscious or unconscious.

opportunity to respond to whatever choice the decision maker makes. In *Priors and Desires* Nature moves first, and while decision makers can make inferences about the true state, nothing they do can change it. As a matter of fact, *Priors and Desires* decision makers are ambiguity neutral, and ambiguity averse ‘pessimists’ (in the [Maccheroni et al. \(2006\)](#) sense) are realists.¹³

2.1 Generalized state-space

The *Priors and Desires* model has so far been presented for the expositionally convenient case of a finite state-space, but practical applications often require an infinite, continuous, or otherwise more complicated environment.

Let Ω denote now any set of states (not necessarily finite), and let Σ denote a σ -algebra of subsets of Ω called “events”. Let F denote the set of all Σ -measurable mappings from Ω to \mathbb{R} . Equation 2 generalizes to the following expression for any event $E \in \Sigma$:

$$p_r(E|s) \propto \int_{\Omega} p_0(\omega|s) e^{\psi r(\omega)} d\omega. \quad (6)$$

2.2 Normal distribution case

The following example is interesting in itself, and is also useful for predicting the behavior of subjects in the experiment.

Example 1. Normal distribution

Suppose $p_0(\cdot)$ has a normal pdf with mean μ and variance σ^2 , and that the stakes (in utility terms) are linear in the state x : $r(x) = ax + b$ for some $a, b \in \mathbb{R}$,¹⁴ then

$$p_r(x) \propto e^{-\frac{(x-\mu)^2}{2\sigma^2}} e^{\psi(as+b)} \propto e^{-\frac{(x-\mu)^2 - 2\psi a\sigma^2 s}{2\sigma^2}} \propto e^{-\frac{(x-(\mu+\psi a\sigma^2))^2}{2\sigma^2}} \quad (7)$$

Hence, p_r also has a normal pdf with variance σ^2 , but with a mean of $\mu + \psi a\sigma^2$. The bias is therefore proportional to both a and to σ^2 . The former represents the reference stakes: the greater a is (in absolute terms) the stronger the dependence of the decision maker’s utility on the state, while the latter represents the degree of uncertainty: the larger σ^2 is, the more

¹³What they have at stake has no impact on the choice they end up making.

¹⁴For example, the stakes in dollar terms are linear in x and the decision maker is risk neutral.

uncertainty there is about the state. Bigger stakes and greater uncertainty result in a bigger bias.

3 Experiment

This section describes the experiment. The implementation and protocol are in Section 3.1, the specifics of the belief elicitation procedure in Section 3.2, and the theoretical predictions of *Priors and Desires* and *Optimal Expectations* in Section 3.3.

3.1 Implementation and protocol

The experiment was conducted at the Center for Experimental Social Science at the University of Oxford. Students registered for participation in an experiment, and were only told that it would require about an hour of their time. Economics, finance, and business students were excluded.¹⁵ Sessions were conducted in the afternoon over a total of six days. There were 12 sessions altogether, half of which were 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. 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.

The experiment consisted of 13 periods, including one training period and 12 earning periods. 12 different price charts (Figure 2) were used for the earning periods, with the order of presentation randomized independently between subjects. At the end of the experiment one earning period was chosen randomly for each subject, and the subject 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. The charts were adapted from

¹⁵The concern was that such students would ignore their subjective judgment in making their predictions of future prices.

real financial markets, and were selected to include a variety of situations.¹⁶ As compared with using simulated data, this was felt to make the task more interesting, and to maximize the chance that subjects use their subjective judgment in making their predictions. 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 or no incentive for speed. 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 (if the day 100 price of wheat is £4,000) and £12,000 (if it's £16,000). 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 3.2. *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 *Farmers* stood to gain from high wheat prices, whereas *Bakers* were better off with low prices.

Sessions differed in the scale of the accuracy bonus and in 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

¹⁶The source was historical stock prices, scaled and shifted to fit into a uniform range.

Table 1: The number of subjects in each condition, defined by a particular combination of maximum accuracy bonus and stakes. There were between 2 and 4 sessions in each condition, each consisting of subjects of only one type.

Condition		
Maximum bonus ^a	stakes ^b	number of subjects
£1	1	49 (25 <i>Farmers</i> , 24 <i>Bakers</i>)
£2	1	26 (13 <i>Farmers</i> , 13 <i>Bakers</i>)
£5	1	44 (23 <i>Farmers</i> , 21 <i>Bakers</i>)
£1	0.5	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*).

notional profit.¹⁷ Table 1 lists the number of subjects in each condition.

3.2 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, 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 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

¹⁷In sessions with lower stakes, subjects received an additional £3, so that the average payoff was the same as in the baseline sessions.

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

$$q(x) = (1 - \epsilon)\mathcal{N}(x|m, (\lambda r)^{-2}) + \epsilon \quad (8)$$

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

$$b(x) = \alpha \log (q(x)/\epsilon) \quad (9)$$

where x is the true day 100 price, and α is a parameter which determines the maximum bonus level.¹⁸ As $q \geq \epsilon$ (Equation 8), 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:

$$\begin{aligned} \mathbb{E}_P[b(x)] &= \int p(x)\alpha \log \frac{q(x)}{\epsilon} dx = \alpha \left(\int p(x) \log \frac{q(x)}{p(x)} dx \right. \\ &\left. + \int p(x) \log p(x) dx - \log \epsilon \right) = \alpha \left(-D_{\text{KL}}(P||Q) - H(P) - \log \epsilon \right) \end{aligned} \quad (10)$$

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

¹⁸The logarithmic scoring rule was introduced in Good (1952). See Gneiting and Raftery (2007) for a recent discussion and comparison to other scoring rules.

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 $Q = P$.¹⁹

The scoring rule works best if subjects are risk neutral and beliefs are well approximated by a density in the family described by Equation 8. 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.²⁰ Consider a risk-averse *Farmer*. Her profit 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 8. 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.

3.3 Predictions

This section develops the predictions of *Priors and Desires* and of the hedonic self-deception model *Optimal Expectations* (Brunnermeier and Parker, 2005). The following timing framework is used: 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 also consume anticipatory utility; at $t = 2$ the day 100 price is revealed, and payoffs are realized. Risk neutrality

¹⁹This 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.

²⁰Blanco et al. (2008) find evidence of hedging in belief reporting when opportunities are transparent and incentives are strong. Armantier and Treich (2010) discuss hedging in probability elicitation.

over small stakes is assumed throughout. It is further assumed that beliefs about the day 100 price can be represented by a distribution from the family described by Equation 8. Given these assumptions predictions reveal beliefs.

3.3.1 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. Payoffs in the experiment are realized at $t = 2$, and consist of two components: profit and 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 9. The $t = 0$ maximand can thus be written as follows:

$$W = \eta\mathbb{E}_Q[\phi\kappa x + b(x)] + (1 - \eta)\mathbb{E}_P[\phi\kappa x + b(x)] + l \quad (11)$$

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 10, we can rewrite Equation 11 as follows:

$$\begin{aligned}
W &= \eta \mathbb{E}_Q[\phi \kappa x + b(x)] + (1 - \eta) \mathbb{E}_P[\phi \kappa x + b(x)] + l \\
&= \eta(\phi \kappa \mu_1 - \alpha H(Q) - \alpha D_{\text{KL}}(Q||Q) - \alpha \log \epsilon) \\
&\quad + (1 - \eta)(\phi \kappa \mu_0 - \alpha D_{\text{KL}}(P||Q) - \alpha H(P) - \alpha \log \epsilon) + l \\
&= \eta(\phi \kappa \mu_1 - \alpha H(Q)) - (1 - \eta)\alpha D_{\text{KL}}(P||Q) + C
\end{aligned} \tag{12}$$

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,²¹ 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 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:

$$\begin{aligned}
\frac{\partial W}{\partial \mu_1} &= \eta \phi \kappa + \eta \frac{\partial H(Q)}{\partial \mu_1} - (1 - \eta) \alpha \frac{\partial D_{\text{KL}}(P||Q)}{\partial \mu_1} \\
&= \eta \phi \kappa - (1 - \eta) \alpha \frac{(\mu_1 - \mu_0)}{\sigma^2}
\end{aligned} \tag{13}$$

²¹That is, the highest possible price of £16,000 if a *Farmer*, and the lowest possible price of £4,000 if a *Baker*.

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

$$\mu_1 - \mu_0 = \phi \left(\frac{\eta}{1 - \eta} \right) \left(\frac{\kappa\sigma^2}{\alpha} \right) \quad (14)$$

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 14 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

$$\mathbb{E}[\mu_1 - \mu_0] = \mathbb{E}[\mu_1] - \mathbb{E}[\mu_0] = \phi \mathbb{E} \left[\frac{\eta}{1 - \eta} \right] \left(\frac{\kappa\sigma^2}{\alpha} \right) \quad (15)$$

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

$$b_{\text{optimal expectations}} = 2\mathbb{E} \left[\frac{\eta}{1 - \eta} \right] \left(\frac{\kappa\sigma^2}{\alpha} \right) \propto \frac{\kappa\sigma^2}{\alpha} \quad (16)$$

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.

3.3.2 Priors and Desires

The *Priors and Desires* model is described in Section 2. Since risk neutrality is assumed, we can use money units to represent the subjects' initial stakes. These correspond to the subject's unconditional payoff: $r(x) = \phi\kappa x$, where x is the day 100 price, κ represents 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 3.3.1, that indifference beliefs are normal: $P = \mathcal{N}(\mu_0, \sigma^2)$. According to Example 1 the biased beliefs are also normal with the same variance, and with a mean shifted in proportion to the coefficient of relative optimism ψ , the slope parameter κ , and the variance σ^2 . In other words, $P_r = \mathcal{N}(\mu_1, \sigma^2)$, where

$$\mu_1 - \mu_0 = \phi\psi\kappa\sigma^2 \tag{17}$$

This equation describes the bias in the beliefs of some particular individual, and is the *Priors and Desires* analogue of Equation 14. By analogy with Section 3.3.1, the expected difference in beliefs between *Farmers* and *Bakers* is

$$b_{\text{priors and desires}} = 2\mathbb{E}[\psi]\kappa\sigma^2 \propto \kappa\sigma^2 \tag{18}$$

Comparing this result to Equation 16, 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 18 is independent of α .

4 Results

This section presents the results of the experiment, starting with a test of the wishful thinking hypothesis, and continuing with the comparative statics of the bias. Parameter estimates and statistical test results are presented in summary form in Tables 2 and 3. The first table presents results using the entire sample, and the second presents the corresponding results with outlier subjects removed. The issue of outliers is discussed in Section 4.1. Figures 4 and 5 provide a graphical illustration of the results in Table 3.

4.1 Wishful thinking

The wishful thinking prediction is that *Farmers* predict higher prices than *Bakers*. Figure 4 shows histograms of the mean prediction reports across all charts. There is a great deal of overlap, and the lowest (highest) prediction is actually made by a *Farmer* (*Baker*). Nevertheless, looking at the histogram

it does seem as if *Farmers* generally predict higher prices. Summary statistics confirm this impression: the mean prediction of approximately 63% of *Farmers* is above the median, and the mirror image of that is true for 62% of *Bakers*. The overall mean in the two groups is £10,118 and £9,728 respectively.

The statistical significance of these observations can be tested using regression analysis. Let y_{ij} denote the prediction made by subject i in chart j , and let d_F denote a dummy for *Farmers*. Given the random allocation we can use the following regression model:

$$y_{ij} = \sum_j (\beta_j d_F + \mu_j) d_j + \delta_i + \epsilon_{ij} \quad (19)$$

where d_j is a dummy for chart j , and δ_i and ϵ_{ij} are the error terms. The μ_j terms represent the population mean prediction of *Bakers* in each of the different charts, and the β_j terms represent the difference in predictions between *Farmers* and *Bakers*. For the purpose of testing for the existence of wishful thinking, it is convenient to take the expectation over j to obtain the following simple regression model:

$$y_i = \beta d_F + \mu + \delta_i \quad (20)$$

where y_i is the mean prediction across charts for subject i , μ is the mean for all *Bakers*, and δ_i is the error term. β represents the expected value of the wishful thinking. The wishful thinking hypothesis is that $\beta > 0$. The OLS estimate of Equation 19 is $\hat{\beta} = 390$, and the null hypothesis that $\beta \leq 0$ is rejected with a p -value of 0.0016.

As is evident from Figure 4, four subjects made predictions that are out of line with all other subjects.²² Outliers can have a disproportionate effect on linear regressions, and it is therefore interesting to repeat the analysis on a sample that excludes these outliers. The revised estimate is $\hat{\beta} = 430$, and the null hypothesis is rejected with a p -value of 0.0001. Thus, the null hypothesis that $\beta \leq 0$ is strongly rejected whether or not we include outliers in the regression.

²²These include the *Farmer* with the lowest predictions, the *Baker* with the highest predictions, and the two *Bakers* with the lowest predictions.

4.2 Cost of holding biased beliefs

Self-deception models predict a decrease in the magnitude of the wishful thinking as a function of the cost of holding biased beliefs, whereas judgment bias models predict no such decrease. Biased beliefs are costly, since they lead to biased predictions, which are likely to be off target. Thus, the greater the bias, the lower is the accuracy bonus that the subject can expect to obtain. The cost of holding biased beliefs is, therefore, an increasing function of the accuracy bonus scale. More specifically, if we are prepared to assume risk neutrality over small stakes, it follows that the cost of holding biased beliefs is a linear function of the accuracy bonus scale. If we further assume that the benefit of biased beliefs is proportional to the subjective expectation of the day 100 price (as is the case, for example, in the Optimal Expectations model) we obtain the testable prediction that the magnitude of the bias should be inversely proportional to the accuracy bonus scale.

The magnitude of the bias was estimated in sessions with a maximum bonus size of £1, £2, and £5 using a generalization of Equation 20 which allows for different levels of bias in different groups of subjects.²³ In addition, the data was used to fit a model which allows for a power function dependence of the bias magnitude on the maximum bonus scale:

$$y_i = 0.5\beta t_i b_i^\gamma + \mu + \delta_i \quad (21)$$

where b_i is the accuracy bonus scale in the session to which subject i was allocated, and other notation is the same as in Equation 20. The value of γ was determined by maximum likelihood estimation, and standard errors were computed using a quadratic approximation to the log likelihood in the vicinity of the maximand. The judgment bias prediction is that $\gamma = 0$, and the self-deception prediction is that $\gamma < 0$. The more specific self-deception prediction obtained under the assumptions discussed above is that $\gamma = -1$.

The results depend on whether outliers are included. The estimated bias for the three sets of sessions 382, 320, and 575 if outliers are excluded, and

²³Let d_k denote a dummy for sessions with maximum bonus k . Replacing β in Equation 20 with $\sum_k \beta_k d_k$ is insufficient, since if the magnitude of the wishful thinking varies between sessions with different bonus size, so would the mean prediction of *Bakers*. However, the effect on the predictions of *Farmers* and *Bakers* should be exactly the same (with an opposite sign), so that the mid-point between the mean prediction in the two groups should not vary with the maximum bonus size. The solution, therefore, is to replace d_F with $d_F - 0.5$, giving μ precisely this interpretation.

206, 461, and 662 if they are retained. The maximum likelihood estimates for the dependence of the bias on the size of the accuracy bonus are, respectively, $\hat{\gamma} = 0.275$ and $\hat{\gamma} = 0.659$. Thus, if anything, the magnitude of the bias seems to an *increasing* function of the bonus scale. These surprising results are not a great fit to either model, but there is a major difference: the prediction of the self-deception model that $\gamma = -1$ is strongly rejected in both regressions (p -values of 0.0139 and 0.0056 respectively), whereas the prediction of the judgment bias model that $\gamma = 0$ is not rejected (p -values of 0.4382 and 0.1254 respectively). Thus it is entirely possible that the true value of γ is zero, and that the increase in the data is due to random noise. This interpretation is particularly convincing if one believes that the regressions with outliers excluded provide a better test than regressions that include the outliers.

It is worth, however, to entertain the possibility that the increase is not merely random noise, so that even in the limit of $N \rightarrow \infty$ we would see an increasing pattern in the data. Such a pattern cannot be explained by any model I am aware of, but it does not seem so strange if we note that spending time trying to predict the day 100 price is costly. Thinking about the day 100 price is necessary for forming an opinion about it, so that only the predictions of subjects who pay attention could possibly be affected by wishful thinking. Assuming some fixed cost for paying attention, it follows that a higher bonus would translate into more subjects choosing to pay attention, and hence a larger bias.²⁴

4.3 Subjective uncertainty

According to both types of model the magnitude of the bias should be increasing in the degree of subjective uncertainty. This prediction can be tested by splitting the 12 charts into two equal sized groups, defined by the degree of subjective uncertainty in the chart, and estimating a regression model, which allows for the magnitude of the wishful thinking to vary between the two groups.²⁵ Two different measures of subjective uncertainty were used.

²⁴Let M denote the magnitude of wishful thinking if a subject pays full attention, and let $\lambda \in [0, 1]$ denote the subject's actual attention. Suppose λ is an increasing as a function of the bonus scale, and that M is independent of the bonus scale (as predicted by the judgment bias model). The actual wishful thinking, λM , would then be an increasing function of the bonus scale.

²⁵It is also necessary to allow for the possibility that the mean prediction is different in the two groups of charts.

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.

The results in Tables 2 and 3 show a much bigger estimated bias in the high uncertainty group of charts, consistent with the predictions of both self-deception and judgment bias models. p -values for the null are between 0.0388 and 0.0589. These results are illustrated graphically in the second and third panels of Figure 5. Panel 2 plots the estimated wishful thinking against the mean prediction confidence in the chart, and panel 3 plots the same data against the within group prediction variance.

4.4 Stakes

The two classes of model also predict that the magnitude of the bias should be increasing in what the subjects have at stake 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).

The magnitude of the bias was estimated separately in these two subsamples using a similar model to that of Section 4.2. The estimated bias with outliers excluded was 257 in the low stakes subsample, and 469 in the high stakes subsample. These results are consistent with the prediction that the magnitude of the bias is linear in the stakes ($p < 0.9323$). Similar results were obtained when outliers were included in the regression. However, the sample size in the low stakes sessions is small, and while results provide an excellent fit to predictions, the null hypothesis that the bias is not any smaller in the low stakes subsample is not rejected ($p < 0.2279$).

4.5 Confidence in the accuracy of predictions

This section seeks to answer the following question: are the predictions of subjects who are confident in their predictions more or less biased than the predictions of subjects who lack confidence in their predictions? The wishful

thinking hypothesis predicts a positive correlation, as long as some subjects are more optimistically biased than others. *Farmers*, for example, gain both from high prices and from accurate predictions. Hence, wishful thinking should cause them to expect higher prices than they would otherwise, and at the same time to be more confident that their predictions are accurate. On the other hand, subjects who don't believe in their ability to predict the day 100 price have less to lose from making biased predictions. Other things being equal, the self-deception model would therefore predict *less* bias in confident subjects. Overall, therefore, self-deception models are ambiguous about the relationship between confidence and the bias level, whereas judgment bias models imply a positive relationship.

Confident subjects were defined by whether their mean reported prediction confidence across the 12 charts was above the median, and, separately, by their answer to a post-experiment questionnaire question asking whether they believe prices in financial markets can be predicted.²⁶ A similar testing methodology to that of Section 4.2 was followed, allowing for a different level of bias in confident and non-confident subjects.

The results in Tables 2 and 3 show a pretty clear positive relationship between bias magnitude and confidence. If outlier subjects are included in the estimate, the test is only statistically significant if the questionnaire question is used to define confident subjects (p -value=0.0698). If, however, outlier subjects are excluded, the test is statistically significant regardless of how confident subjects are defined (p -values of 0.0695 and 0.0413 respectively). This positive correlation result fits the prediction of judgment bias models.

5 Discussion

The paper started with the observation that wishful thinking has powerful implications to decision making, but that while we have plenty of evidence suggesting that wishful thinking is real and pervasive, there remains ample room for doubt.

²⁶The 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.

The theoretical contribution of the paper is the *Priors and Desires* model of wishful thinking. Instead of modeling wishful thinking as a choice (as do models of self-deception), in *Priors and Desires* wishful thinking is a judgment bias. A key implication of this difference is that *Priors and Desires* is consistent with a pervasive wishful thinking bias that affects any and all decisions involving subjective judgment of likelihood, including high-stakes decisions where wishful thinking can be potentially very costly to the decision maker.

The empirical contribution of the paper consisted of an experiment that provides a simple test of wishful thinking, and makes it possible to study its comparative statics. Despite incentives for hedging, subjects gaining from high prices predicted systematically higher prices than subjects gaining from low prices. This result is readily explained as a consequence of wishful thinking bias, and unlike studies in more complicated environments, is difficult to explain otherwise.

The experiment offers no obvious opportunities for subjects to deceive themselves, and its results are therefore hard to explain as a consequence of self-deception. By increasing the size of the accuracy bonus it was possible to make it more costly for subjects to bias their beliefs, but no decrease in the magnitude of the bias was observed. This result is particularly hard to account for in a (hedonic) self-deception model, and further suggests that the wishful thinking bias in the experiment is better understood as a consequence of a judgment bias. Other comparative statics results were also consistent with a judgment bias model, such as *Priors and Desires*.

The implication is that the process people use to make subjective judgments of likelihood is prone to wishful thinking bias. This does not imply that it is the *only* source of wishful thinking bias,²⁷ but it does suggest that we should expect to see wishful thinking whenever it is predicted by a model, such as *Priors and Desires*. Of course, *Priors and Desires* implies that *all* subjective judgments of likelihood are affected by wishful thinking.²⁸ The conclusion, therefore, is that wishful thinking is indeed real and pervasive, and that it is something to keep in mind whenever we model decisions that depend on subjective judgments of likelihood.

²⁷In particular, it certainly does not imply that self-deception is not real, and that we should not expect a potentially larger wishful thinking bias in circumstances where self-deception implies a bias.

²⁸Unless, of course, the decision maker happens to be completely indifferent between all states.

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, *Journal of Economic Perspectives* **11**(1): 109–126.
- Benabou, R. and Tirole, J. (2002). Self-confidence and personal motivation, *Quarterly Journal of Economics* **117**(3): 871–915.
- 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 Lovo, 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.
- Cowgill, B., Wolfers, J., Wharton, U. and Zitzewitz, E. (2009). Using Prediction Markets to Track Information Flows: Evidence from Google, Mimeo.

- Festinger, L. and Carlsmith, J. (1959). Cognitive consequences of forced compliance, *Journal of Abnormal Psychology* **58**(2): 203–10.
- 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.
- Hoffman, M. (2011a). Learning, Persistent Overconfidence, and the Impact of Training Contracts, Mimeo.
- Hoffman, M. (2011b). Overconfidence at Work and the Evolution of Beliefs: Evidence from a Field Experiment, Mimeo.
- 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.
- Maccheroni, F., Marinacci, M. and Rustichini, A. (2006). Ambiguity Aversion, Robustness, and the Variational Representation of Preferences, *Econometrica* **74**(6): 1447–1498.
- Malmendier, U. and Tate, G. (2008). Who makes acquisitions? CEO overconfidence and the market’s reaction, *Journal of Financial Economics* **89**(1): 20–43.
- 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.

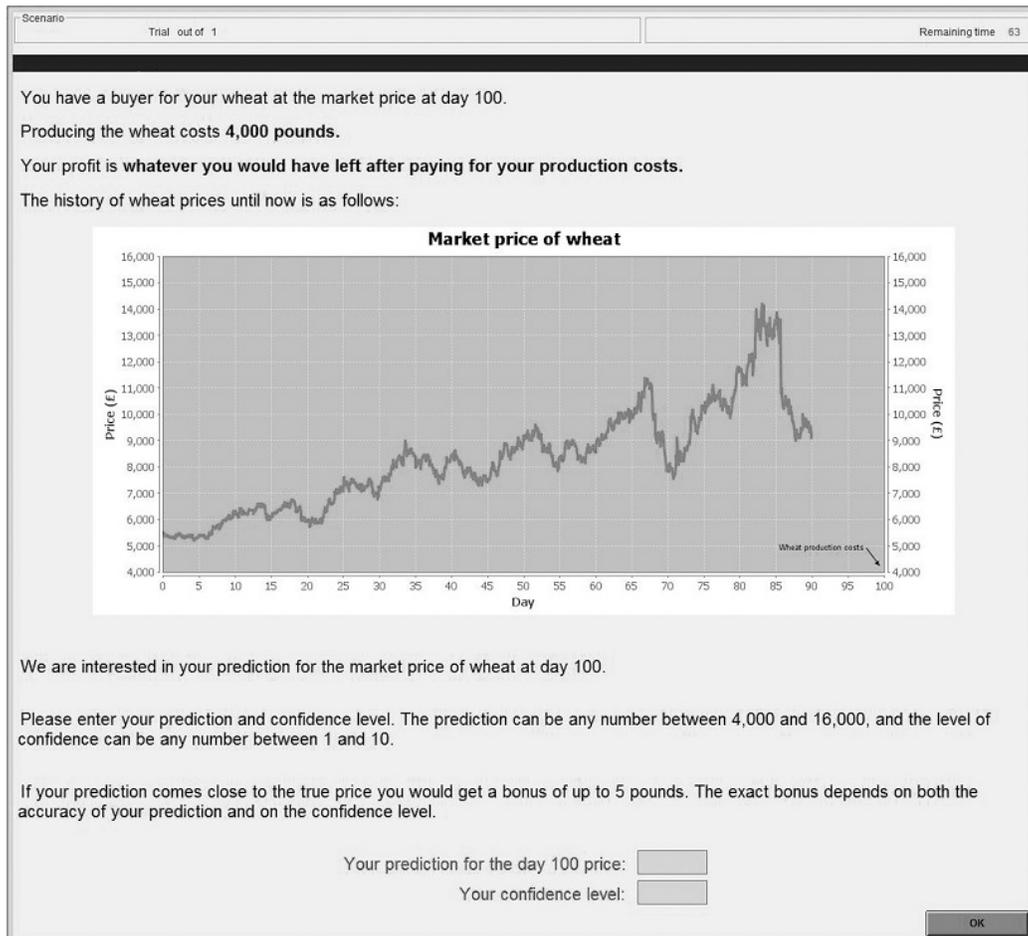


Figure 1: The interface of the *Farmers* treatment with a maximum accuracy bonus of £5. The interface of the *Bakers* treatment was similar, except: (a) the first three lines were: “You have a buyer for £16,000 worth of bread from your bakery. At day 100 you will get the money from the order, and will 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”.

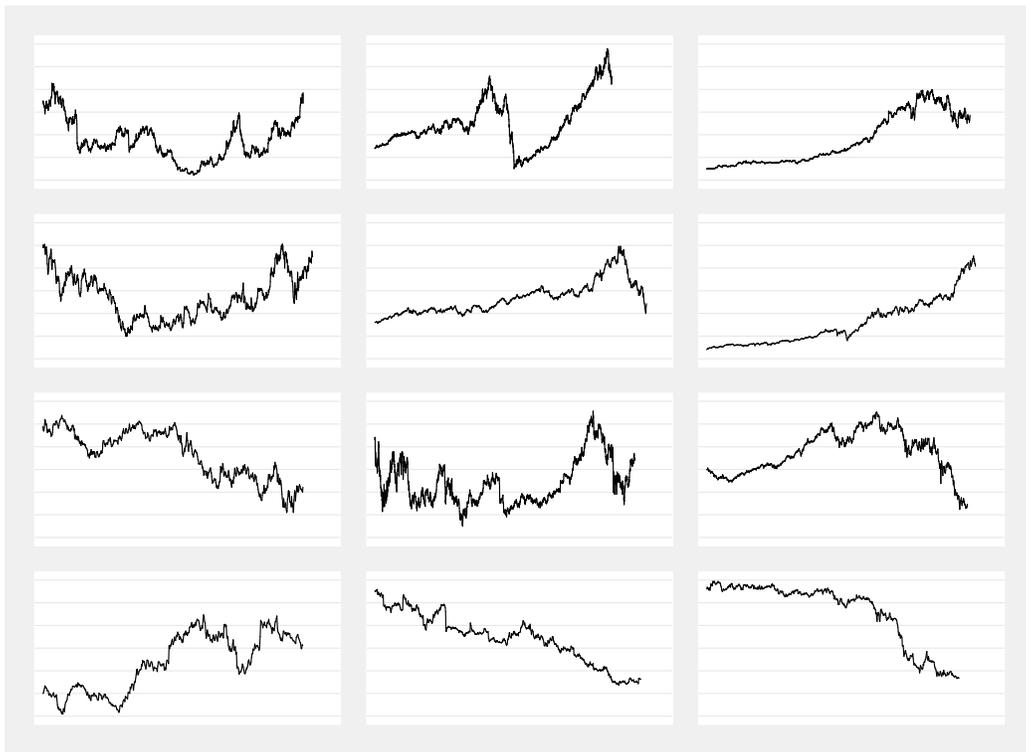


Figure 2: The charts used in the 12 earning periods. The x -axis represents time, ranging from day 0 to day 100, and the y -axis represents price, ranging from £4,000 to £16,000. The data for the charts were adapted from historical equity price data, shifted and scaled to fit into a uniform range. Figure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

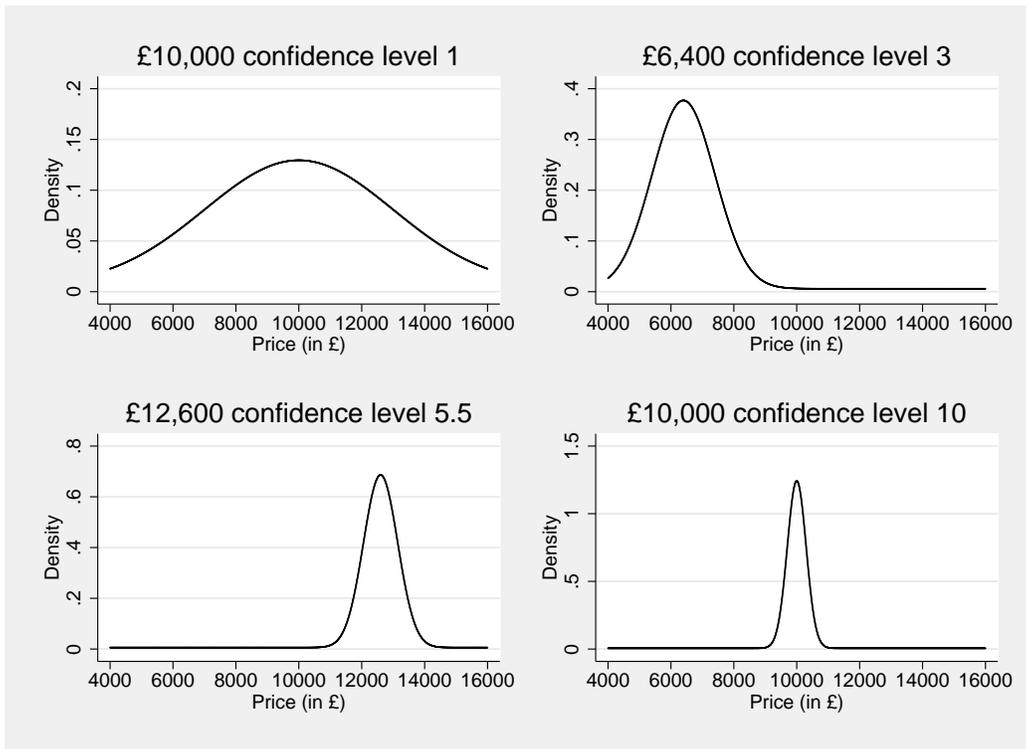


Figure 3: The examples of distributions used in the instructions. Each distribution is characterized by a prediction and a confidence level. These examples were used in explaining the prediction elicitation procedure. They were particularly useful in establishing a reference for the 1-10 scale that was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

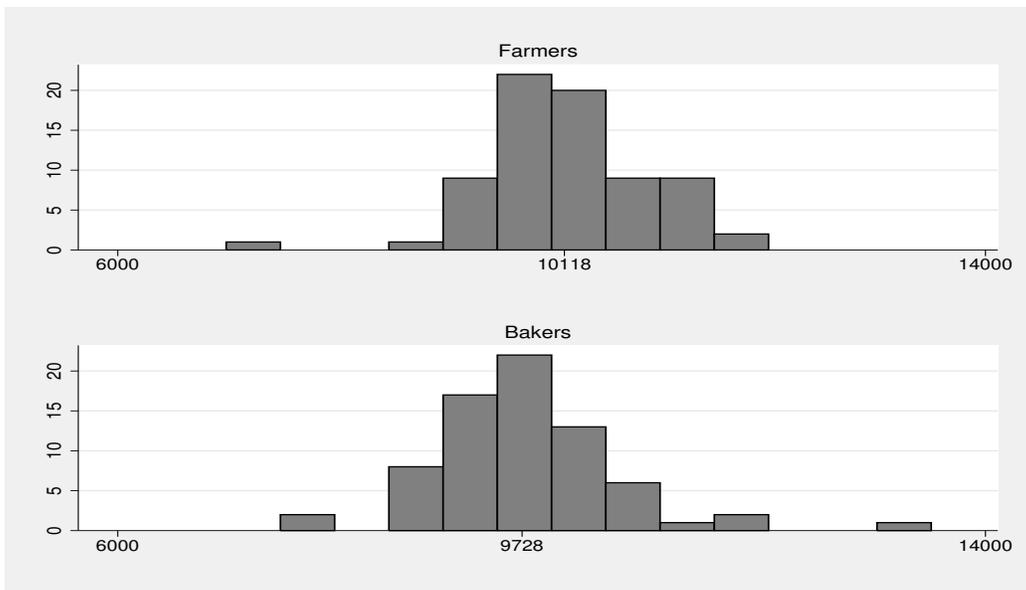


Figure 4: Histogram—split by subject type—of the mean prediction made by all subjects. The mean prediction in the two groups was 10118 and 9728 respectively.

Table 2: Optimism bias and its comparative statics with outlier subjects included (see Section 4.1). Table 3 shows the same regressions with outliers excluded.

	Sample	Estimated bias ^a	Observations ^b
	All subjects	390*** (s.e. 130)	145
	<i>negative ?</i>	$p < 0.0016$	
Cost of holding biased beliefs	Accuracy bonus: <i>low</i> (£1)	206 (s.e. 181)	75
	Accuracy bonus: <i>medium</i> (£2)	461* (s.e. 307)	26
	Accuracy bonus: <i>high</i> (£5)	662*** (s.e. 236)	44
	ML exponent ^c	0.659 (s.e. 0.465)	145
	<i>exponent = 0 ?</i>	$p < 0.1254$	
	<i>exponent = -1 ?</i>	$p < 0.0056$	
Subjective uncertainty	Chart uncertainty: <i>low</i>	200* (s.e. 136)	145
	Chart uncertainty: <i>high</i>	579*** (s.e. 165)	145
	<i>low > high ?</i>	$p < 0.0388$	
	Within chart variance: <i>low</i>	204** (s.e. 120)	145
	Within chart variance: <i>high</i>	576*** (s.e. 179)	145
	<i>low > high ?</i>	$p < 0.0423$	
Stakes in the day 100 price	Stakes: <i>low</i>	254 (s.e. 308)	26
	Stakes: <i>high</i>	420*** (s.e. 144)	119
	<i>high = 2 · low ?</i>	$p < 0.8902$	
	<i>low ≥ high ?</i>	$p < 0.3132$	
Confidence in ability to predict prices	Average confidence: <i>low</i>	314** (s.e. 188)	70
	Average confidence: <i>high</i>	461*** (s.e. 182)	75
	<i>low > high ?</i>	$p < 0.2884$	
	Prices predictable? <i>no</i>	201 (s.e. 182)	74
	Prices predictable? <i>yes</i>	587*** (s.e. 185)	71
	<i>no > yes ?</i>	$p < 0.0698$	

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

^b Each observation is the mean prediction of a given subject across all 12 charts, except in the part, where charts are split into two groups of 6 charts each.

^c A power function was fitted for the dependence of the bias on the accuracy bonus. p-values based on LR test. The standard error is the standard deviation of a 2nd order approximation of the likelihood function around the maximand.

Table 3: Optimism bias and its comparative statics with outlier subjects excluded (see Section 4.1). Table 2 shows the same regressions with outliers included.

	Sample	Estimated bias ^a	Observations ^b
	All subjects	430*** (s.e. 110)	141
	<i>negative ?</i>	$p < 0.0001$	
Cost of holding biased beliefs	Accuracy bonus: <i>low</i> (£1)	382*** (s.e. 153)	73
	Accuracy bonus: <i>medium</i> (£2)	320 (s.e. 262)	25
	Accuracy bonus: <i>high</i> (£5)	575*** (s.e. 200)	43
	ML exponent ^c	0.275 (s.e. 0.352)	141
	<i>exponent = 0 ?</i>	$p < 0.4382$	
	<i>exponent = -1 ?</i>	$p < 0.0139$	
Subjective uncertainty	Chart uncertainty: <i>low</i>	282*** (s.e. 108)	141
	Chart uncertainty: <i>high</i>	578*** (s.e. 155)	141
	<i>low > high ?</i>	$p < 0.0589$	
	Within chart variance: <i>low</i>	261*** (s.e. 100)	141
	Within chart variance: <i>high</i>	598*** (s.e. 164)	141
	<i>low > high ?</i>	$p < 0.0400$	
Stakes in the day 100 price	Stakes: <i>low</i>	257 (s.e. 256)	26
	Stakes: <i>high</i>	469*** (s.e. 122)	115
	<i>high = 2 · low ?</i>	$p < 0.9323$	
	<i>low ≥ high ?</i>	$p < 0.2279$	
Confidence in ability to predict prices	Average confidence: <i>low</i>	264** (s.e. 156)	69
	Average confidence: <i>high</i>	589*** (s.e. 153)	72
	<i>low > high ?</i>	$p < 0.0695$	
	Prices predictable? <i>no</i>	243* (s.e. 153)	72
	Prices predictable? <i>yes</i>	624*** (s.e. 156)	69
	<i>no > yes ?</i>	$p < 0.0413$	

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

^b Each observation is the mean prediction of a given subject across all 12 charts, except in the “subjective uncertainty” part, where charts are split into two groups of 6 charts each.

^c A power function was fitted for the dependence of the bias on the accuracy bonus. p-values based on LR test. The standard error is the standard deviation of a 2nd order approximation of the likelihood function around the maximand.

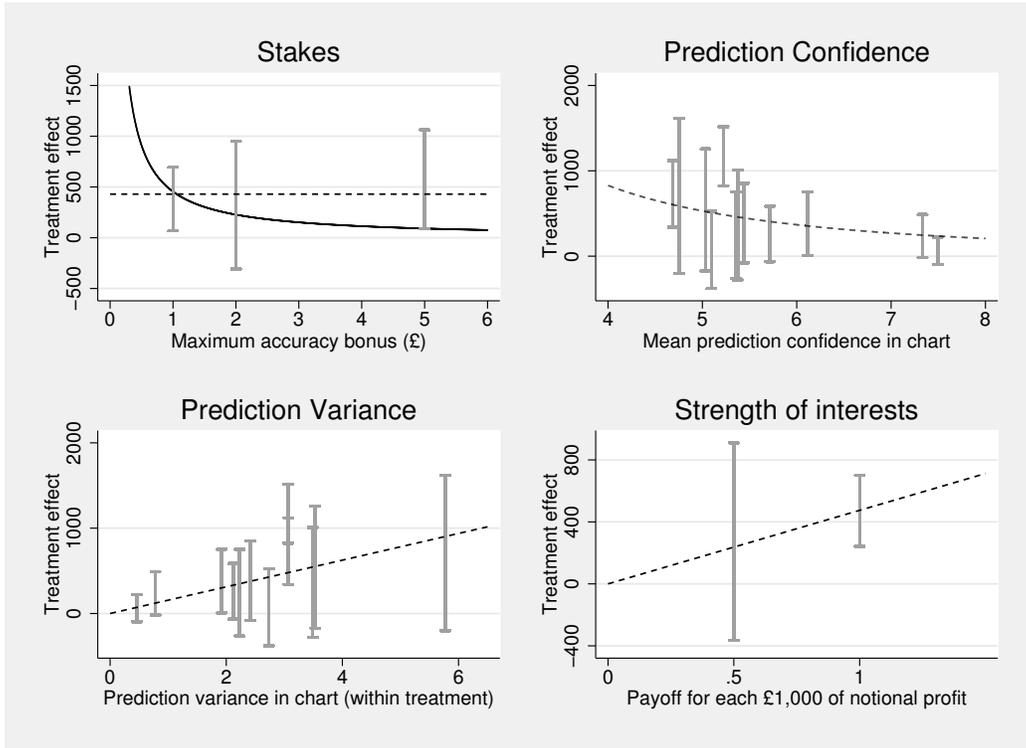


Figure 5: The comparative statics of the wishful thinking. The panels show a 95 percent confidence interval for the difference in predictions between *Farmers* and *Bakers* in different subsamples, with outlier subjects excluded. The first panel focuses on the cost of holding biased beliefs, as represented 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.