

Testing Models of Belief Bias: An Experiment

Alexander Coutts[‡]
New York University

November 20 2014

[Click for latest version.](#)

Abstract

Optimistic beliefs affect important areas of economic decision making, yet direct knowledge on how belief biases operate remains limited. To better understand these biases I conduct an experiment examining beliefs about binary events with financial stakes. By varying financial prizes in outcomes, as well as incentive payments for accuracy, the experiment is able to distinguish between two leading theories of optimistic belief formation that differ in their assumptions about how such beliefs are constrained. The Optimal Expectations theory of Brunnermeier and Parker (2005) models beliefs as being constrained through the future costs of holding incorrect beliefs, while the Affective Decision Making model of Bracha and Brown (2012) argues that beliefs are constrained by mental costs of distorting reality. The results suggest that people hold optimistically biased beliefs, and comparative statics indicate that these beliefs are not constrained by increasing the costs of making inaccurate judgments. In fact, the results support the theory of Bracha and Brown (2012), as observed bias is *increasing* in the size of incentive payments for accuracy.

JEL classification: C91, D03, D80, D81, D83, D84.

Keywords: Beliefs · Optimism · Pessimism · Overconfidence · Anticipation · Affective expected utility

*19 West 4th Street FL 6, New York, NY 10012. alexander.coutts@nyu.edu

[‡]Acknowledgements: This research has been generously supported by a grant from the Russell Sage Foundation. I am heavily indebted to my advisor David Cesarini for numerous discussions and comments. I am grateful for helpful comments from Hunt Allcott, Andrew Demers, Guillaume Fréchet, Nicole Hildebrandt, David Low, Amnon Maltz, Joseph Mullins, Giorgia Romagnoli, Andrew Schotter, Emilia Soldani, Tobias Salz, Séverine Toussaert, Christopher Woolnough, Sevgi Yuksel, as well as numerous participants at ESA International and North American meetings in 2014.

1 Introduction

Optimistic beliefs play an important role in decision making, yet a lack of direct evidence hinders the ability of researchers to model these biases. Accurately modeling optimism is critical for developing theory and informing policy. Beliefs affect decisions such as saving for retirement, starting a new business, or investing in the stock market.¹ This has motivated theorists to challenge the benchmark rational model of decision making under uncertainty. While substantial theoretical progress has been made, rigorous tests of existing theory and direct evidence about optimism are scarce.² The contribution of this paper is to provide this direct evidence, and distinguish between two leading ways to model optimism.

Early thinkers such as Bentham (1789) discussed the pleasures associated with holding positive expectations about the future. Contemporary theories of optimism have incorporated this idea that individuals get utility from anticipation of the future. Akerlof and Dickens (1982) implicitly use this framework in their model of Cognitive Dissonance when they consider the benefits of workers having reduced fear about the probability of an accident at the workplace. There must also be costs to holding optimistic beliefs, as experience suggests people do not always hold extreme beliefs. In contrast to the relative consensus regarding what the benefits of optimism are, there is disagreement in theory over how these costs should be modeled.

Two approaches have been taken in the literature to model the costs of holding optimistic beliefs. First, holding distorted beliefs may lead an individual to make worse decisions. This approach is taken by Brunnermeier and Parker (2005) (henceforth BP): Optimal Expectations trade off the anticipatory benefits of optimism with the costs of worse decision making. Here optimal beliefs are *not* directly constrained by reality. In contrast, the second approach is to model a direct cost of distorting reality. As individual optimism strays further from the “truth”, it becomes increasingly costly to internally justify holding these beliefs. Bracha and Brown (2012) (henceforth BB) model a strategic game between an emotional process and a rational process, that emphasizes these direct costs of distorting reality.

To illustrate the difference between the two approaches, consider the example from

¹Evidence that optimistic beliefs play a role in these situations can be found in Puri and Robinson (2007), Landier and Thesmar (2009), and Easterwood and Nutt (1999) respectively.

²Within psychology, there exist some direct experimental tests of optimism, however participants in these experiments are typically not provided financial incentives for accurate responses, and these studies are often not designed with the aim of distinguishing or testing theory. See for example Vosgerau (2010).

Oster et al. (2013). They have panel data on patients at risk of Huntington Disease (HD), including both objective risk status, as well as subjective risk perceptions. HD is a degenerative neurological disorder with onset around age 40, that significantly reduces life expectancy. They find that individuals are optimistic about their risk of having HD and subsequently are less likely to be tested for the disease. Patterns in the data are not consistent with a standard rational model, but are consistent with the optimal expectations model.³ According to BP, the costs of these beliefs are living with the disease, yet not knowing and making uninformed life decisions. These costs are weighed against the benefits, determining the optimal level of bias. On the other hand, in BB, as an individual becomes more and more optimistic about his risk status for HD, his beliefs are stretched further and further from reality. In determining the optimal belief, it becomes harder and harder for the individual to believe he is of very low risk, e.g. in face of evidence from doctors who are telling him he is at high risk for HD.

Both BB and BP model benefits to holding optimistically biased beliefs, however the policy prescriptions for de-biasing individuals are very different. The model of BP argues that the only way to reduce optimism bias is by increasing the indirect costs of worse decision making. The BB model is suggestive that information campaigns that raise the costs of distorting reality have the potential to be effective. In the case of HD, the model of BP implies that forcing high risk individuals to confront their risk status will not affect testing rates, whereas in BB this will increase the direct costs of distorting beliefs, leading to higher testing rates.

This paper distinguishes between these two models by eliciting beliefs of individuals about binary events in which they have a financial stake. The elicitation procedure induces truthful reporting of beliefs for individuals, regardless of whether they are standard Rational Expectations (RE) agents, BP agents, or BB agents. Individuals are given an income stake in an event, and must report the probability this event occurs. The two comparative statics involve varying the level of income, and varying the incentive payment for accurate reporting. The elicitation procedure follows the BDM method applied to lotteries, which is attractive for two key reasons. First, the procedure is incentive compatible for any risk preferences, and second, it theoretically generates opposite comparative static predictions for the two models of interest. Increasing the income stake exacerbates opti-

³Oster et al. (2013) also provide evidence that other models of belief bias such as the information aversion model of Koszegi (2003) do not fit the patterns in the data. It is worth emphasizing, while they do not consider the model of Bracha and Brown (2012), it does fit the patterns in the data.

mism bias in both BP and BB models, which is unsurprising as both model the benefits to optimism similarly. However, increasing the incentive payment for accuracy results in a lower bias for BP agents, but a *higher* bias for BB agents.

To preview the intuition for this result, in BB, increasing incentive payments does not alter the costs of distorting reality, but affects anticipatory benefits as expected payoffs in the experiment increase. In BP, the cost of making worse decisions is precisely the loss in expected payment from reporting a biased belief, which dominates any anticipatory benefits.

Finally, as a qualitative test of the predictions of BB, I examine probability reports across different domains. Two events involve a random process that is “objective”, in the sense that there exists consensus that baseline probabilities exist and can be calculated.⁴ The events differ in how cognitively demanding it is to determine the baseline probability. The other two events do not have obvious baseline probabilities that can be calculated, and may depend on subject ability. BB make the qualitative prediction that as events become less objective, costs of distortion will change, while BP predict no change in the costs of distortion, and thus no change in the degree of bias.

To preview the results, I find evidence of optimistic beliefs across all domains and treatments. Giving subjects a large financial stake in an event leads to an increase in bias, consistent with both models. I also find that giving subjects larger incentive payments for accurate reports leads to an *increase* in optimism bias, a prediction made uniquely by the model of BB. BP make a somewhat more intuitive, but opposite prediction that larger incentive payments for accuracy will lead to less bias, which is rejected by the data. Also consistent with BB, is that, qualitatively, the degree of bias is increasing as the random process that determines events become less objective.⁵

The next section outlines related literature and experimental evidence of optimism in psychology and economics. I then introduce the experiment in the form of a stylized model and derive the theoretical comparative static results outlined above. I next discuss the experimental design, followed by results, and a concluding discussion.

⁴I am not aware of a rigorous definition of objective in this context. One can see Gilboa and Schmeidler (2001) for a definition, which agrees with the characterization in this paper.

⁵This statement of course is not directly testable, as I do not have a rigorous definition of the objectivity of a random process.

2 Related Literature

2.1 Experimental

There are a number of experiments testing or seeking to understand models of belief bias, often in the realm of overconfidence. Such models are related in the sense that they assume individuals benefit from holding positive views about their self-image or ability, compared with models of optimism that assume benefits to holding positive views about the future.⁶ There are a number of situations where overconfidence and optimism overlap, for example the prospect of getting a future salary increase may be a signal of high ability that also is a direct increase in income.

In most experiments on overconfidence, the focus is not only on prior beliefs, but also on how subjects update their beliefs about personal qualities such as intelligence. These experiments are motivated by models of biased information seeking, and/or biased information processing. Some evidence that individuals might overweight positive information is found in studies by Mobius et al. (2013) and Eil and Rao (2011).⁷ Both papers also find a positive relationship between seeking information about ability when such information is likely to confirm prior beliefs.

In contrast to experiments that examine overconfidence, Mayraz (2013) examines beliefs over outcomes and finds some evidence of optimistic beliefs in a novel experiment where sessions were divided into “farmers” and “bakers” and income depends on a hypothetical price of wheat. While the motivation is most similar to this paper, in the current experiment I use an incentive compatible elicitation procedure with no opportunities for hedging. Additionally I ensure that randomization of the primary treatment is at the individual level, ensuring subjects’ earning expectations are invariant to the treatment.⁸ Further, the experiment of this paper maintains a close connection to theoretical work on optimism bias, which will lay the groundwork for the design and interpretation of results.

⁶Models of overconfidence include Benabou and Tirole (2002) and Köszegi (2006) among others.

⁷There are studies that do not find such asymmetries in updating, for example Grossman and Owens (2012) finds no such pattern (however they look at absolute rather than relative performance), while Ertac (2011) actually finds the opposite pattern.

⁸A design concern is that if subjects have earnings expectations for their time in the lab, they can back out moments of the distribution of a random event, and this can generate beliefs that appear optimistic. In Section 4 I explain how I overcome this potential confound.

2.2 Theoretical

For most theoretical work on optimism bias, the benefits from holding optimistic beliefs are that they lead to gains from anticipation of the future.⁹ Akerlof and Dickens (1982) was one of the earliest models to motivate such beliefs in economics, describing the decision of a worker to choose (incorrectly) to believe that the probability of an accident in his workplace is low.¹⁰ Loewenstein (1987) explicitly outlined a model of anticipatory benefits to consumption, theorizing that this could explain why individuals might optimally delay consumption.

This idea that individuals might get utility from anticipation was further explored by Caplin and Leahy (2001) who modeled choice behavior in such a world, as well as Landier (2000) who introduced a model where anticipatory feelings were nurtured through a biased information process.¹¹ A more recent model similar in spirit to Landier (2000) is found in Mayraz (2013). In his Priors and Desires model, individuals bias perceived probabilities of various states proportional to the relative attractiveness of the states. Increasing payoffs in a given state increases the optimistic bias, while there are no costs associated with optimism in the model.¹²

Most relevant are the two key theories of this experiment, those of Brunnermeier and Parker (2005) (BP) and Bracha and Brown (2012) (BB). The next section describes these theories in detail.¹³

⁹Other models of belief bias are related including models of biased memory, biased information seeking, or other benefits to ego-preserving beliefs, such as motivation or social signalling, as described in Benabou and Tirole (2002). In the experiment, there is no memory component and the structure of receiving information is exogenous. Additionally the treatment parameters are given exogenously, and not related to personal qualities.

¹⁰Specifically, there are psychological costs of fear increasing in the probability an accident will occur. Here, anticipation is over losses, and hence there are anticipatory benefits to believing such states are less likely.

¹¹Related is Koszegi (2010) who modeled an individual optimizing simultaneously beliefs and behavior, defining a personal equilibrium concept.

¹²Formally the model is agnostic over whether individuals are optimistic, realist (rational), or pessimistic.

¹³A large literature exists examining the theory of decisions under uncertainty/ambiguity. An entry level overview can be found in Etner et al. (2012). Bracha and Brown (2012) provide an axiomatic characterization of their model, showing that an alternative interpretation of the model is that it captures ambiguity seeking behavior. In particular the model is identical to the variational preferences model of Maccheroni et al. (2006) (MMR), when Axiom A.5 (uncertainty aversion) is replaced with uncertainty seeking. The difference in attitudes towards uncertainty/ambiguity may arise from differences in source: in BB the ambiguity is endogenous, in MMR it is exogenous.

3 Theory

In this section I summarize two different theories of biased beliefs that generate different, testable predictions in the experiment. Next, I will describe the experiment, in the form of a stylized model, and discuss implications for the two theories. The final task is to characterize optimal beliefs under each of these models and summarize the relevant comparative statics that distinguish the two theories from each other, and from standard rational expectations (RE).

3.1 Summary of Brunnermeier and Parker (2005) (BP)

In the optimal expectations framework of BP, the key tradeoff is between the anticipatory benefits agents get from holding optimistic beliefs and the costs of these beliefs due to worse decision-making. I consider a two period model with consumption occurring only in the second period, an identical setting to the stylized model I will introduce in Section 3.3.¹⁴

The first period involves individuals choosing an action that only affects second period consumption.¹⁵ Utility is given by $u(\cdot)$, an increasing, twice differentiable function that is independent of time. There are a finite number of n states $S \equiv \{s_1, \dots, s_n\}$ with state s_j occurring with objective probability π_j . Consumption in state s_j is given by c_j . Denote $\hat{\pi}_j$ as the subjective probability an individual assigns to state s_j .

The agent behaves as if he were a standard expected utility maximizer, but facing subjective probabilities $\{\hat{\pi}\}$ ¹⁶ rather than objective probabilities RE agents are assumed to face, $\{\pi\}$. In particular, the agent solves the following optimization problem, choosing a vector of actions \mathbf{x} , where $g_j(\mathbf{x})$ is a mapping from actions to state specific consumption.

The agent makes choices in period 1, but does not consume until period 2.

¹⁴BP additionally allow a role for memory, where individuals may retrospectively gain utility by remembering past consumption. When there are only two periods and consumption only occurs in the second period, memory plays no role in well-being.

¹⁵This setup is identical to the Portfolio Choice example presented by BP Section II. They include a discount factor β but it does not affect the decision problem, so I omit it in the present exercise.

¹⁶I use the notation $\{\hat{\pi}\}$ to represent the $n \times 1$ vector of probabilities $\{\hat{\pi}_1, \dots, \hat{\pi}_n\}$.

$$\begin{aligned} \max_{\mathbf{x}} \quad & \sum_{s_j \in S} \hat{\pi}_j \cdot u(c_j) \\ \text{s.t.} \quad & c_j = g_j(\mathbf{x}) \quad \forall j \end{aligned} \tag{1}$$

The optimal solution to this problem is $\mathbf{x}^*(\{\hat{\pi}\})$, which leads to consumption in each state, $c_j^*(\{\hat{\pi}\})$. This gives optimal consumption in every state, given probabilities $\{\hat{\pi}\}$. Thus far, the individual is a standard expected utility maximizer given optimal beliefs: $\{\hat{\pi}\}$.

I now turn to the fundamental component of the BP model, the selection of these optimal beliefs, $\{\hat{\pi}\}$. Total welfare is a weighted average of anticipation given beliefs in the first period, and actual consumption in the second period. To avoid degenerate beliefs, I follow Oster et al. (2013), Spiegel (2008)¹⁷, and Bridet and Schwardmann (2014), and allow for the possibility that utility from anticipation is lower than utility from actual consumption.¹⁸ Mathematically, I introduce a degree of anticipation parameter $\gamma \in [0, 1]$. $\gamma = 1$ is the modeling assumption in BP, while lower levels of γ correspond to lower benefits to optimism, with $\gamma = 0$ representing a benchmark RE agent.

Given optimal consumption from Equation 1 above, utility from anticipation in the first period is:

$$\gamma \cdot \sum_{s_j \in S} \hat{\pi}_j \cdot u(c_j^*(\{\hat{\pi}\})) \tag{2}$$

In the second time period utility is actual consumption, which occurs with the objective probabilities $\{\pi\}$:

¹⁷Spiegel (2008) similarly nest BP as a special case of weighting anticipation by α and consumption by $1 - \alpha$. He shows that more generally there exists α^* such that beliefs are degenerate. In many applications when the costs of holding extreme beliefs is not negative and large in absolute value, $\alpha^* \leq \frac{1}{2}$.

¹⁸That is, I allow for the possibility that today, the thought of consuming a pizza tomorrow (measured today) gives me a lower amount of pleasure than what I get from actually eating the pizza tomorrow (measured tomorrow).

$$\sum_{s_j \in S} \pi_j \cdot u(c_j^*(\{\hat{\pi}\})) \quad (3)$$

Optimal expectations are those that maximize the time average of lifetime utility, subject to satisfying standard laws of probability.¹⁹ Here this is the average of anticipation in the first period and actual consumption in the second period. Substituting in optimal consumption (as a function of beliefs) $c_j^*(\{\hat{\pi}\})$, optimal expectations, $\{\hat{\pi}\}$, are the solution to the maximization problem:

$$\begin{aligned} \max_{\{\hat{\pi}\}} \quad & \frac{1}{2} \cdot \gamma \cdot \sum_{s_j \in S} \hat{\pi}_j \cdot u(c_j^*(\{\hat{\pi}\})) + \frac{1}{2} \cdot \sum_{s_j \in S} \pi_j \cdot u(c_j^*(\{\hat{\pi}\})) \quad (4) \\ \text{subject to} \quad & \text{(i) } \sum_{j=1}^n \hat{\pi}_j = 1, \text{ (ii) } \hat{\pi}_j \geq 0, \text{ and (iv) } \hat{\pi}_j = 0 \text{ if } \pi_j = 0 \end{aligned}$$

Note that these subjective expectations are chosen once and only once. In a multi-period model this is equivalent to an agent being endowed with optimal priors, and using Bayes' rule to update beliefs over time.

Looking at Equation 4, setting $\gamma = 0$ eliminates all benefits to holding biased beliefs. The solution is $\{\hat{\pi}\} = \{\pi\}$, i.e. the standard beliefs of an RE agent.

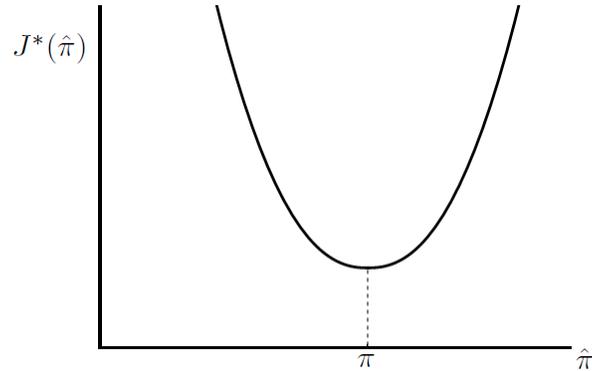
3.2 Summary of Bracha and Brown (2012) (BB)

In this model, choice and beliefs are determined simultaneously by the outcome of an intrapersonal game between two cognitive processes, a rational and an emotional process. The key difference between the models of BB and BP are the constraints on optimistic beliefs. In BB a mental cost function $J^*(\{\hat{\pi}\})$ constrains beliefs for the emotional process, whereas in BP it is the costly decision errors that result from holding biased beliefs.

I maintain the same setup as in the earlier section, in particular the same utility function $u(\cdot)$ and n states, $s_j \in S$, with π_j the objective probability of state s_j . The emotional process chooses a subjective probability vector $\{\hat{\pi}\}$ that maximizes the following, which

¹⁹In particular: (i) $\sum_{j=1}^n \hat{\pi}_j = 1$ (ii) $\hat{\pi}_j \geq 0$ for all j and (iv) $\hat{\pi}_j = 0$ if $\pi_j = 0$. Note the last assumption in BP means that individuals cannot assign positive subjective probability to zero probability events.

Figure 1: Mental Cost Function



Two-state example: here π is the true probability of one of the states, with $1 - \pi$ being the probability of the other state, $\hat{\pi}$ is the distorted probability.

takes choices by the rational process \mathbf{x} , and consequently c_j as given.²⁰

$$\max_{\{\hat{\pi}\}} \sum_{s_j \in S} \hat{\pi}_j \cdot u(c_j) - J^*(\{\hat{\pi}\}) \quad (5)$$

$J^*(\{\hat{\pi}\})$ is a function of Legendre-type: strictly convex, essentially smooth function on the interior of the probability simplex Δ .²¹ The function reaches a minimum at $\{\hat{\pi}\} = \{\pi\}$, and is such that in the limit as $\hat{\pi}_j$ goes to either 0 or 1 (for any $\pi_j \in (0, 1)$), $J^*(\{\hat{\pi}\})$ approaches infinity at a higher rate than the utility function, guaranteeing that holding extreme beliefs is never optimal. An example of the two state case can be seen in Figure 1.

The intuition for such a cost function is based on evidence from the psychology literature, that people use mental strategies such as biased search to justify their beliefs. As desired beliefs are further away from the “truth” search costs to support these beliefs become greater.

The rational process maximizes expected utility given beliefs $\{\hat{\pi}\}$, which is identical

²⁰In BP optimal beliefs are chosen under the assumption that the agent makes choices after being endowed with these beliefs. Here by contrast, beliefs and choices are simultaneously chosen.

²¹For more details, see Bracha and Brown (2012).

to Equation 1 in the summary of the BP model in the previous section. Note that unlike the BP model, the rational process does not observe beliefs, as they are determined simultaneously by the intrapersonal game.

$$\begin{aligned} \max_{\mathbf{x}} \quad & \sum_{s_j \in S} \hat{\pi}_j \cdot u(c_j) \\ \text{s.t.} \quad & c_j = g_j(\mathbf{x}) \quad \forall j \end{aligned}$$

Each of these two payoff functions determine the best response functions. The intersection of the best response functions are the Nash equilibria of the game. BB show that there are an odd number of locally unique pure strategy Nash equilibria.

3.3 Theoretical Predictions

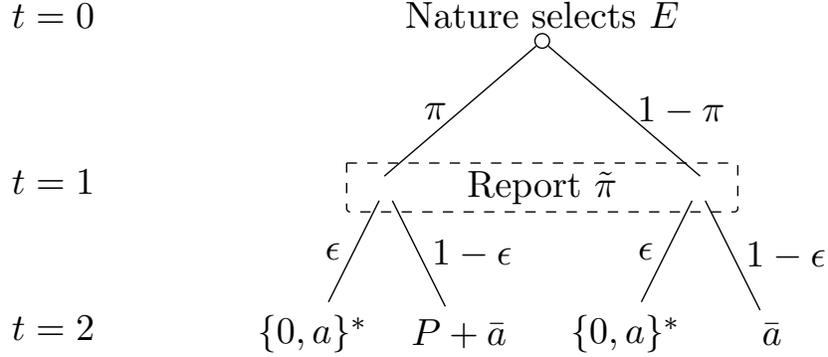
I now outline a stylized model of the experiment and later will characterize optimal behavior for the three types of agents: BP, BB, and the benchmark RE. There is a binary event of interest, E which occurs with true probability $\pi \in (0, 1)$ or does not occur with probability $1 - \pi$. First nature determines whether E occurs, next an agent is asked to report a probability $\tilde{\pi}$ that event E occurs, critically having no information about the occurrence of E .²² $\tilde{\pi}$ corresponds to action \mathbf{x} in the previous section.

Payoffs are determined in the following way, which is described in Figure 2. In order to ensure that agents do not wish to hedge their probability reports, the world is partitioned into two disjoint states, the accuracy state and the prize state.²³ With probability $\epsilon \in (0, 1)$ the individual is paid solely according to her reported belief $\tilde{\pi}$ about whether event E occurred using the incentive compatible lottery method to elicit beliefs with an accuracy payment of $a > 0$ (*accuracy state*).

²²The reason that outcomes are determined before agents give probability reports is to maintain consistency with a secondary component of the experiment that involved updating, given signals. Providing signals required that these outcomes were known in advance.

²³Hedging will be present whenever utility is not linear, for example with a concave utility function an individual would prefer to report a lower than truthful $\tilde{\pi}$, since this will smooth consumption over the different states of the world. Karni and Safra (1995) show that without this partition, no elicitation procedure exists that induces truthful reporting, a fact that is sometimes overlooked in the experimental literature, see Armantier and Treich (2013).

Figure 2: Timing of Stylized Model



*In the accuracy state the payoff is either 0 or a , depending on the reported belief $\tilde{\pi}$ and whether E occurred, according to the lottery method.

In the second state occurring with probability $1 - \epsilon$, the individual receives a guaranteed payment $\bar{a} \geq a^{24}$ and receives an additional $P \geq 0$ if E occurs, but receives nothing extra if E does not occur (*prize state*). Her report of $\tilde{\pi}$ is no longer relevant in this prize stage.

The lottery method is incentive compatible even for non-risk neutral agents²⁵, and is implemented as follows. The individual is asked to submit a report $\tilde{\pi}$, the probability that event E is realized. A random number r is drawn from any distribution with full support on $[0, 1]$, here I use the uniform distribution. If the individual's estimate $\tilde{\pi} \geq r$ she is paid an amount $a > 0$ if E occurs, and 0 if E does not occur. If $\tilde{\pi} < r$, she plays a lottery that pays out a with probability r , and 0 otherwise. It is relatively simple to show that truthful reporting dominates any other report by the agent.²⁶

In addition to the lottery method having the benefit of being incentive compatible, unlike other proper-scoring rules such as the quadratic scoring rule, the method also has advantages in terms of the theoretical predictions. Later in this section I will generate comparative static predictions for both BP and BB models. Using the lottery method allows

²⁴The payment of \bar{a} is to ensure that the prize state is always preferred to the accuracy state.

²⁵See Karni (2009) for a more detailed description of lottery method, though the method itself has been described in a number of earlier papers.

²⁶Note that a number of other methods commonly employed in the literature are not incentive compatible, including proper scoring rules such as the quadratic scoring rule. See Armantier and Treich (2013) for theory and experimental evidence of the performance of some of these rules.

one to unambiguously distinguish the two models using the comparative static predictions, a result that is not true of proper scoring rules in general.

The optimal probability report for the benchmark RE agent, $\tilde{\pi}$ can be solved algebraically for this stylized model:²⁷

$$\max_{\tilde{\pi}} \left\{ \epsilon \cdot \left[\tilde{\pi} \cdot \left(\pi \cdot u(a) + (1 - \pi) \cdot u(0) \right) + (1 - \tilde{\pi}) \cdot \left(\frac{1 + \tilde{\pi}}{2} \cdot u(a) + \left(1 - \frac{1 + \tilde{\pi}}{2} \right) \cdot u(0) \right) \right] + (1 - \epsilon) \cdot \left[\pi \cdot u(P + \bar{a}) + (1 - \pi) \cdot u(\bar{a}) \right] \right\} \quad (6)$$

Note that P does not factor in to the optimization since $\tilde{\pi}$ has no bearing on this state of the world.²⁸ Hence I arrive at $\tilde{\pi}^{RE} = \pi$ for RE agents. I now turn to the question of incentive compatibility for BP or BB agents. Before directly solving these agents' problems, I must impose more structure on the scope of what random elements of the experiment might be subject to distorted beliefs.

An important assumption maintained throughout this paper is that individuals only distort the primary probability of interest: the outcome of the binary event E occurring with baseline probability π . In particular, the experimental design ensured that the determination of secondary random elements of the experiment were highly transparent, including the nature of how the accuracy versus the prize stage is determined, and subjects had experience with the lottery method and uniform distribution.²⁹

Having said this, in Appendix A I take steps to show how the theoretical predictions might be affected if subjects distort the probability of being in the accuracy state (with baseline probability ϵ), or distort perceptions of the distribution for how r (in the lottery method) is drawn. There I introduce sufficient conditions such that all the theoretical results in this paper continue to hold, even permitting this extension. Given the assumption that only the probability of the outcome of E is distorted, I arrive at Proposition 1.

Proposition 1. *An agent in the experiment truthfully reports her belief, regardless of whether she is a BP, BB, or RE agent.*

²⁷For $r \sim U[0, 1]$, the probability $r < \tilde{\pi} = F(\tilde{\pi}) = \tilde{\pi}$ and $E[r|r > \tilde{\pi}] = \frac{\int_{\tilde{\pi}}^1 r dF(r)}{1 - F(\tilde{\pi})} = \frac{1 + \tilde{\pi}}{2}$.

²⁸This result assumes independence between the prize and accuracy states.

²⁹A literal interpretation of Brunnermeier and Parker (2005) would lead an individual to optimally believe a coin flip with payoff of one dollar if heads and ninety-nine cents if tails, that heads would occur with absolute certainty. I argue this is not the context that the model is meant to apply to. Rather, individuals must face events that have some element of subjectivity in order for any bias to emerge.

Proof. Denote the belief of a BP agent by $\hat{\pi}^{BP}$, the belief of a BB agent by $\hat{\pi}^{BB}$, and a standard RE agents by $\hat{\pi}^{RE} = \pi$. Since $\epsilon > 0$ the report $\tilde{\pi}$ is relevant to all agents, while it has already been established that the lottery method is incentive compatible, using the uniform distribution. In BP, agents choose actions as standard EU maximizers, given optimal beliefs. In BB, the rational process also makes decisions as a EU maximizer, given beliefs from the emotional process. Thus the objective problem for BP agents and the rational process of a BB agent will look nearly identical to Equation 6, only now accounting for distorted beliefs $\hat{\pi} \in \{\hat{\pi}^{BP}, \hat{\pi}^{BB}\}$ rather than the truth π .

$$\max_{\tilde{\pi}} \left\{ \epsilon \cdot \left[\tilde{\pi} \cdot \left(\hat{\pi} \cdot u(a) + (1 - \hat{\pi}) \cdot u(0) \right) + (1 - \tilde{\pi}) \cdot \left(\frac{1 + \tilde{\pi}}{2} \cdot u(a) + \left(1 - \frac{1 + \tilde{\pi}}{2} \right) \cdot u(0) \right) \right] + (1 - \epsilon) \cdot \left[\hat{\pi} \cdot u(P + \bar{a}) + (1 - \hat{\pi}) \cdot u(\bar{a}) \right] \right\} \quad (7)$$

Setting the resulting first order condition to zero yields:

$$\epsilon \cdot \left[(u(x) - u(0)) \cdot (\hat{\pi} - \tilde{\pi}^*) \right] = 0$$

Leading to the optimal report $\tilde{\pi}^* = \hat{\pi}$. □

Thus, regardless of whether an agent is a BP, BB, or RE type, she will truthfully report her belief $\hat{\pi}$. The next section characterizes what optimal beliefs are for BP agents, followed by BB agents.

3.3.1 Optimal Beliefs in Brunnermeier and Parker (2005) (BP)

In the BP model, utility at time $t = 1$ will be solely the anticipation of second period income realizations (i.e. over earning $P + \bar{a}$ if E occurs, \bar{a} otherwise) and anticipation over the elicitation payment (i.e. over potentially earning a). Recall $\gamma \in [0, 1]$ is utility from anticipation, with $\gamma = 0$ corresponding to a standard RE agent. Utility at time $t = 1$ is thus:

$$\gamma \left[\epsilon \cdot \left(\tilde{\pi} \cdot (\hat{\pi} \cdot u(a) + (1 - \hat{\pi}) \cdot u(0)) + (1 - \tilde{\pi}) \cdot \left(\frac{1 + \tilde{\pi}}{2} \cdot u(a) + \left(1 - \frac{1 + \tilde{\pi}}{2}\right) \cdot u(0) \right) \right) + (1 - \epsilon) \cdot \left(\hat{\pi} \cdot u(P + \bar{a}) + (1 - \hat{\pi}) \cdot u(\bar{a}) \right) \right] \quad (8)$$

Utility at time $t = 2$ from second period consumption depends on the true probability, π :

$$\epsilon \cdot \left(\tilde{\pi} \cdot (\pi \cdot u(a) + (1 - \pi) \cdot u(0)) + (1 - \tilde{\pi}) \cdot \left(\frac{1 + \tilde{\pi}}{2} \cdot u(a) + \left(1 - \frac{1 + \tilde{\pi}}{2}\right) \cdot u(0) \right) \right) + (1 - \epsilon) \cdot \left(\pi \cdot u(P + \bar{a}) + (1 - \pi) \cdot u(\bar{a}) \right) \quad (9)$$

From Proposition 1 it is known that $\tilde{\pi} = \hat{\pi}$. Substituting this value in, the optimal choice of $\hat{\pi}$ is the maximization of 8 + 9, and is solved by:

$$\hat{\pi}^{BP} = \begin{cases} \pi & \text{if } \gamma = 0. \\ \min \left\{ \frac{\pi}{1-\gamma} + \frac{(1-\epsilon)\gamma}{\epsilon(1-\gamma)} \cdot \frac{u(P+\bar{a})-u(\bar{a})}{u(a)-u(0)}, 1 \right\} & \text{if } 0 < \gamma < 1. \\ 1 & \text{if } \gamma = 1. \end{cases} \quad (10)$$

Note that beliefs are restricted to the interval $[0, 1]$, from Equation 10 it is clear that this restriction may bind, leading to the optimal belief $\hat{\pi}^{BP} = 1$. Optimal beliefs are increasing in γ , the weight on anticipatory utility. When $\gamma = 0$ the utility from anticipation disappears, and optimal beliefs coincide with those of a standard RE agent, $\hat{\pi}^{BP} = \pi$.

3.3.2 Optimal Beliefs in Bracha and Brown (2012) (BB)

Since E is binary, the mental cost function can be written as $J^*(\hat{\pi}, \pi)$. It satisfies the following properties: $\lim_{\hat{\pi} \rightarrow 0} |DJ^*(\hat{\pi}, \pi)| = \lim_{\hat{\pi} \rightarrow 1} |DJ^*(\hat{\pi}, \pi)| = +\infty$ and $\lim_{\hat{\pi} \rightarrow 0} J^*(\hat{\pi}, \pi) = \lim_{\hat{\pi} \rightarrow 1} J^*(\hat{\pi}, \pi) = +\infty$. As described in an earlier section, these properties and the assumption that the mental cost function approaches infinity at a higher rate than the utility function as $\hat{\pi} \rightarrow 0$ or $\hat{\pi} \rightarrow 1$ ensure that optimal beliefs are always in the interior, $\hat{\pi}^{BB} \in (0, 1)$.

From Proposition 1, the rational process of a BB agent will truthfully report $\tilde{\pi} = \hat{\pi}$. I thus proceed to the determination of optimal beliefs $\hat{\pi}^{BB}$. Optimal beliefs will be determined by the intersection of the best response functions of the emotional and rational processes.

$$\begin{aligned} \max_{\tilde{\pi}} \left\{ \epsilon \cdot \left[\tilde{\pi} \cdot \left(\hat{\pi} \cdot u(a) + (1 - \hat{\pi}) \cdot u(0) \right) + \right. \right. \\ \left. \left. (1 - \tilde{\pi}) \cdot \left(\frac{1 + \tilde{\pi}}{2} \cdot u(a) + \left(1 - \frac{1 + \tilde{\pi}}{2} \right) \cdot u(0) \right) \right] + \right. \\ \left. (1 - \epsilon) \cdot \left[\hat{\pi} \cdot u(P + \bar{a}) + (1 - \hat{\pi}) \cdot u(\bar{a}) \right] - J^*(\hat{\pi}, \pi) \right\} \end{aligned} \quad (11)$$

Setting the first order condition equal to zero gives the best response function of the emotional process, given $\tilde{\pi}$ from the rational process, (where $J^*(\hat{\pi}, \pi)$ is the first derivative of $J^*(\hat{\pi}, \pi)$):

$$\epsilon \cdot \tilde{\pi} \cdot \left(u(a) - u(0) \right) + (1 - \epsilon) \cdot \left(u(P + \bar{a}) - u(\bar{a}) \right) - J^{*'}(\hat{\pi}, \pi) = 0$$

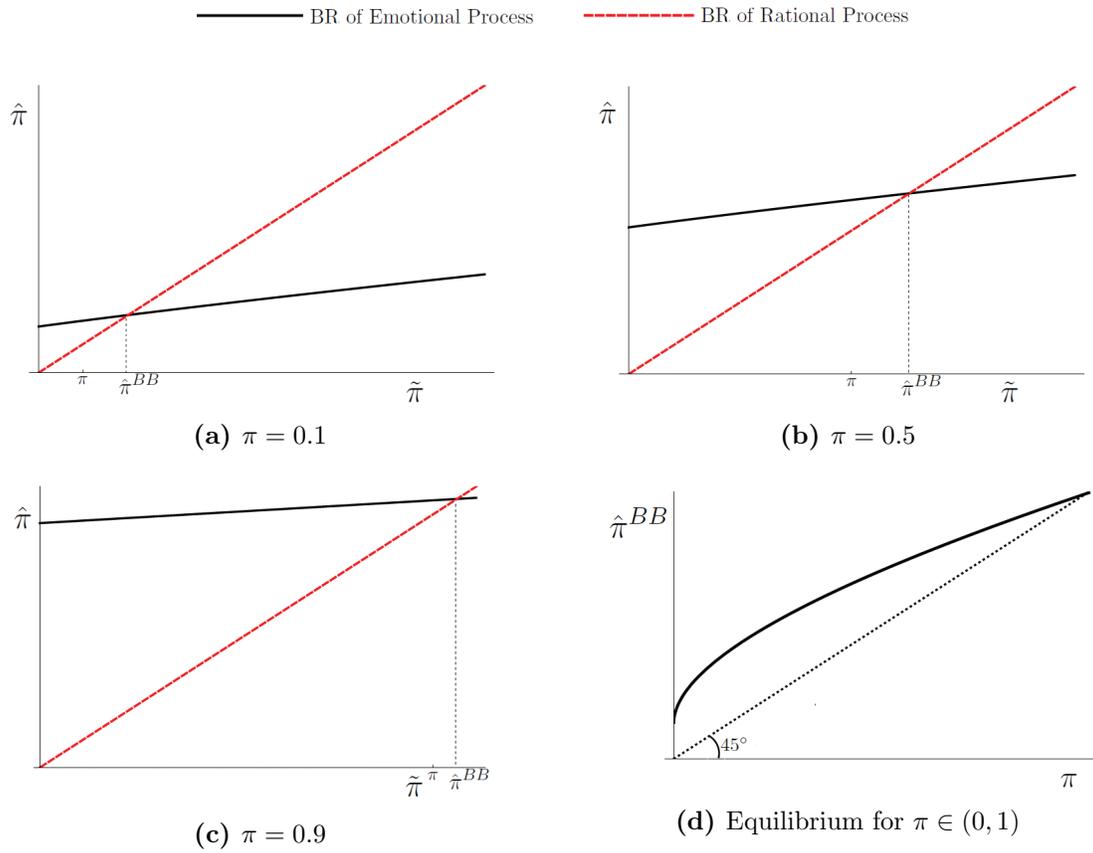
The Nash Equilibrium is the intersection of the best response function of the emotional and rational process. As has been shown, the best response function of the rational process is the truthful report $\tilde{\pi} = \hat{\pi}$. Substituting this into the equation above implicitly defines optimal beliefs $\hat{\pi}^{BB}$, which can be seen in Figure 3 and mathematically in Equation (12).

$$\boxed{\epsilon \cdot \hat{\pi}^{BB} \cdot \left(u(a) - u(0) \right) + (1 - \epsilon) \cdot \left(u(P + \bar{a}) - u(\bar{a}) \right) - J^{*'}(\hat{\pi}^{BB}, \pi) = 0} \quad (12)$$

The two leftmost terms in the equation relate to the benefits of belief bias, while the final term represents the marginal (mental) cost of these beliefs. Regardless of whether this is modeled as a simultaneous moves game or a sequential moves game (with the emotional process as the first mover) the equilibrium is identical.³⁰

³⁰One can show that in Equation 11 if $\tilde{\pi}$ is replaced with $\hat{\pi}$ (the emotional process uses backward induction to determine the rational process' best response), the FOC is identical. The equilibrium would not in general be identical if the rational process moved first.

Figure 3: Best Response (BR) Functions



(a) to (c) illustrate best response curves for $\pi \in \{0.1, 0.5, 0.9\}$, where $\tilde{\pi}$ is the report of the rational process (best response in red) and $\hat{\pi}$ is the belief of the emotional process (best response in black). (d) plots optimal beliefs as a function of the true probability $\pi \in (0, 1)$. The dotted line is the 45° line.

The optimal belief will be $\hat{\pi}^{BB} \geq \pi$. $J^*(\hat{\pi}, \pi)$ takes a minimum at $\hat{\pi} = \pi$, and is decreasing on $[0, \pi]$ and increasing on $[\pi, 1]$. Looking at Equation (12) one can see that $\hat{\pi} = \pi$ gives a strictly higher payoff than any $\hat{\pi} \in [0, \pi)$, since $a > 0$, $P > 0$ and $J^*(\hat{\pi}, \pi) \geq J^*(\pi, \pi)$ for any $\hat{\pi} \in [0, 1]$. The optimum will thus lie in the interval $[\pi, 1]$.

$J^*(\hat{\pi}, \pi)$ is continuous, non-negative on $[\pi, 1]$, equal to zero at $\hat{\pi} = \pi$ and $\lim_{\hat{\pi} \rightarrow 1} J^*(\hat{\pi}, \pi) = +\infty$. Thus a solution to the above equation exists. Additionally, because $J^*(\hat{\pi}, \pi)$ is strictly convex, the solution is unique.

The basis of the experiment is changing the parameters P and a which have different implications for the two models. I now look at comparative static results for the BB agent. First I look at the change in optimal beliefs with respect to a change in the prize P (where $J^{*''}(\hat{\pi}^{BB}, \pi)$ is the second derivative of $J^*(\hat{\pi}, \pi)$). I perturb Equation (12) around the equilibrium to examine how optimal beliefs change as a function of P .

$$\frac{d\hat{\pi}^{BB}}{dP} = \frac{(1 - \epsilon) \cdot u'(P + \bar{a})}{J^{*''}(\hat{\pi}^{BB}, \pi) - \epsilon \cdot (u(a) - u(0))} > 0$$

The denominator of this expression $J^{*''}(\hat{\pi}^{BB}, \pi) - \epsilon \cdot (u(a) - u(0))$ is positive which follows directly from the second order condition for being at a maximum. Next it is possible to look at how optimal beliefs change with respect to the accuracy payment a :

$$\frac{d\hat{\pi}^{BB}}{da} = \frac{\epsilon \cdot \hat{\pi}^{BB} \cdot u'(a)}{J^{*''}(\hat{\pi}^{BB}, \pi) - \epsilon \cdot (u(a) - u(0))} > 0$$

$\frac{d\hat{\pi}^{BB}}{da} > 0$ may seem counter-intuitive, however it is logical upon reflection. The key is that the payoffs a and P have no effect on the marginal costs of holding biased beliefs (costs are only present in the mental cost function), they only affect the marginal benefits. Under the lottery method of belief elicitation, the probability of receiving payment a is increasing in the probability of the event being predicted. Increasing a leads to an increase in the marginal benefit of increasing $\hat{\pi}$, but leaves marginal costs unchanged. This unambiguously increases the BB agent's optimal belief.

Hypotheses

Equations 10 and (12) are the key equations that inform the experimental design. In particular, implicit differentiation gives clear comparative static results that as the financial stake P increases, the degree of bias of $\hat{\pi}$ increases under both models. However, as the accuracy incentive payment a increases, the degree of bias decreases for BP agents, but increases for BB agents.

It has already been shown that for all three agents, BP, BB, and RE, the optimal $\tilde{\pi}^*$ will be to truthfully report optimal beliefs $\hat{\pi}^*$. Given the comparative static predictions for BP and BB models, it is possible to vary the parameters P and a in the lab in order to test the following experimental predictions:³¹

Hypothesis 1:

$$\begin{aligned}\frac{d\hat{\pi}^*}{dP} &> 0 \text{ for BP and BB agents,} \\ \frac{d\hat{\pi}^*}{dP} &= 0 \text{ for a RE agent.}\end{aligned}$$

Hypothesis 2:

$$\begin{aligned}\frac{d\hat{\pi}^*}{da} &< 0 \text{ for an BP agent when } P > 0, \\ \frac{d\hat{\pi}^*}{da} &= 0 \text{ for a RE agent and an BP agent when } P = 0, \text{ and} \\ \frac{d\hat{\pi}^*}{da} &> 0 \text{ for a BB agent.}\end{aligned}$$

4 Experimental Design

This experiment was designed to closely fit the stylized model presented in the previous section that tests the predictions of the Brunnermeier and Parker (2005) (BP) and Bracha and Brown (2012) (BB) models. The motivation and design with testable hypotheses was

³¹In BP, these comparative static predictions require that individuals are *not* at a corner solution for the level of belief bias. This assumption is easily tested, by checking whether or not subjects' probability reports are at the boundary $\hat{\pi} = 1$.

outlined in a grant proposal to the Russell Sage Foundation, prior to the experimental data collection.³² As the primary outcomes of interest in the experiment are elicited beliefs, I utilize the design introduced in Section 3.3 that eliminates incentives for non risk-neutral subjects to hedge payoffs between the accuracy and prize states.³³

The experiment utilizes a 3x2 between subjects design, which can be seen in Figure 4. Three accuracy payment levels, a , were randomized at the session level, low (\$3), moderate (\$10), or high (\$20).³⁴ The accuracy payments were only relevant when subjects ended up in the Accuracy state, revealed at the end of the experiment.

Two prize payment levels, $P \in \{\$0, \$80\}$, were randomized at the subject level. These potential payments were only relevant when subjects ended up in the Prize state. Half of subjects had the chance to earn an extra \$80 if the event occurred, while the other half would earn nothing (\$0) extra. A fixed payment of $\bar{a} = \$20$ was provided conditional on ending up in the prize state, but independent of the event. $\bar{a} \geq a$ was chosen so that the prize state would always be strictly preferred to the accuracy state.³⁵

Rather than only face one event in the lab, subjects faced a sequence of four independent events presented in random order.³⁶ One of the four events was randomly selected at the end of the experiment for payment.

Figure 5 summarizes the four events that all individuals faced. Two events I consider “objective” (using the earlier definition), and involved rolls of the dice that differed in how cognitively demanding it was to calculate the underlying probability. However the outcome of these events was determined by chance, and individuals could not affect these outcomes.

In the other two events, subjects could affect the probability of the outcome. These two events involved performance on both a skill testing quiz, and a scenario where subjects had to guess the weather on a random day in New York City in the previous calendar year.

³²Specifically, details of this design were provided in Grant Proposal #98-14-06 for the Russell Sage Foundation, which outlined that the design would test both BP and BB models.

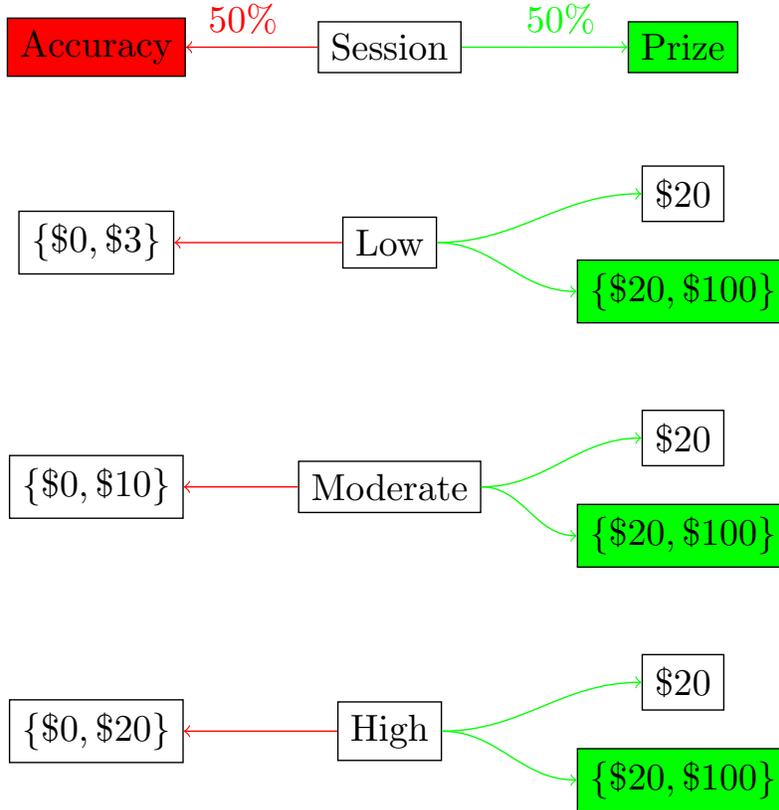
³³This design has been previously utilized with similar aims in Blanco et al. (2010), who also showed that when incentives to hedge are transparent individuals in experiments do take advantage of hedging opportunities.

³⁴Sessions were evenly split across the three treatments. Optimal cost-benefit sample ratio calculations suggested over-sampling low payment relative to high, however this optimal number of additional subjects was less than the size of an average session.

³⁵For any individual with non-degenerate beliefs.

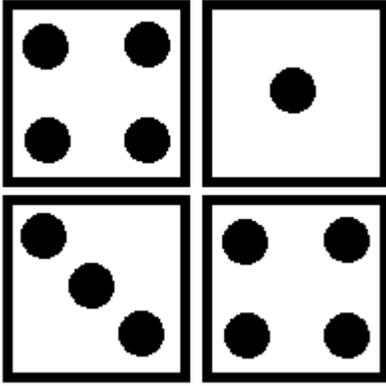
³⁶One of the events, intended as a control was always the final event. The other three events were randomly ordered at the session level.

Figure 4: Experimental Design

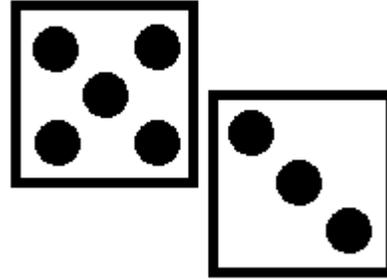


Sessions were allocated between Low, Moderate, and High accuracy payments (the '3' in the '3x2' design). Within sessions, subjects had a 50% chance of ending up in the Prize state or the Accuracy state. In the Accuracy state ex-ante payments were fixed. In the prize state, half of the subjects could potentially earn an extra \$80 if the event occurred, while the other half would receive no such bonus (the '2' in the '3x2' design).

Figure 5: Description of Events



(a) Hard Dice: The event occurred whenever the computer rolled four dice, exactly two out of those four dice was a specified number (e.g. 4). The probability of this is $\binom{4}{2} \left(\frac{1}{6}\right)^2 \left(\frac{5}{6}\right)^2 = \frac{150}{1296}$ or approximately 11.57%.



(b) Easy Dice: The event whenever the computer rolled two dice, two different specified numbers were the only numbers to come up (e.g. 5-3, or 3-5, 3-3, 5-5). The probability of this is $\frac{4}{36}$ or approximately 11.11%.



(c) Weather: The event occurred whenever the individual correctly estimated the average temperature on a specified random day in NYC in the previous year (2013), +/- 5 deg F. In the sample, 25.77% of subjects were in the correct range.

$$x^{3/2} = 64$$

(d) Quiz: The event occurred whenever the individual scored in the top 15% on a skill-testing multiple choice quiz, relative to students in pilot sessions. For a subset of participants the event pertained to a random partner's performance instead of their own.

Both the weather and quiz exercises were completed before the experiment began, hence subjects did not know their potential prize payment P at the time they completed these tasks.

All events had binary outcomes, and the outcome of each event was always determined before subjects submitted their probability reports. Of course, subjects did not know the outcome when they submitted these reports.

The experiment also examined whether there were any differences in beliefs when individuals were given control over choosing the numbers for the two dice events that determined whether the event would occur. 50% of the subjects in each session had to make a choice about the numbers they wanted to see appear for the objective probability events involving dice.³⁷ The motivation for this treatment was to test the hypothesis that individuals given control would be more optimistic about the event occurring, based on psychological evidence about the “illusion of control”, as in Langer (1975).

The third and fourth events involved subject activities that took place before the main elicitation part of the experiment began. The third event was whether a subject scored in the top 15% on a 5 minute skill testing quiz³⁸ that was taken by all subjects. Subjects were incentivized by being truthfully informed that achieving a high score on the quiz would result in an increased chance at earning an extra \$80.³⁹ A random subset of students (30%) were used as a control group where the event occurred based on the performance of a random anonymous partner in the room, rather than their own performance.

The fourth event involved whether subjects correctly estimated the average temperature on a given (random) day in the previous calendar year (2013) in New York City.⁴⁰ As in the quiz question, subjects were given this question before the experiment began, and were told a correct answer would lead to an increased chance of earning \$80.

³⁷The control condition was randomized within subjects, not within events. Hence a subject either had control over both dice events, or none.

³⁸The quiz was a multiple choice quiz consisting of math and verbal questions. To determine whether a subject was in the top 15% they were compared to a reference group of students taking the same quiz during pilot sessions.

³⁹Subject feedback indicated that this was a strong incentive to put in effort on the quiz. Additionally, out of 219 subjects for which I have choice time data, the fastest person finished in 3 minutes (177 seconds). 93% (203) of subjects finished in at least 4 minutes out of the 5 minutes maximum.

⁴⁰Subjects needed to be within a 5 degree Fahrenheit window in order to be correct.

4.1 Timing and Procedures

This experiment was conducted at New York University (NYU), in New York, at the Center for Experimental and Social Science (CESS). Recruitment was done via the CESS online system, where undergraduates are notified by email when the experiment was scheduled, and had the opportunity to sign up. A total of 318 subjects participated⁴¹, in 32 different sessions for an average of 10 subjects per session. The average subject payment was \$24.96 for approximately 75 minutes. All subjects were given a \$10 show up payment in addition to anything earned during the experiment. Due to the nature of the experimental design, final subject payments including show up fee ranged from as low as \$10, to as high as \$110.

In order to address concerns about the difficulty of understanding the lottery method, approximately half (35 minutes) of the experiment consisted of lengthy instructions, along with practice sessions on the computer (in Z-Tree, Fischbacher (2007)) to help subjects get used to the elicitation procedure.⁴²

After the detailed practice which went over the lottery method as well as the payment structure, subjects took the five minute skill testing quiz, followed by answering a question about the weather in NYC on a random day in the previous year, 2013.⁴³ The reason the quiz and weather questions were placed before elicitation was that it was important that subjects did not know what stake they would have in an event when answering these questions. Evidently if some subjects knew they had a chance at earning \$80 for a top performance, while others knew they had no chance at earning \$80, exerted effort could be different.

After the quiz and weather questions, subjects were introduced to their potential stake in the event, $P \in \{\$0, \$80\}$, each equally likely. To ensure this was as transparent as possible I came around with a bag that was filled with a number of poker chips equal to the number of subjects in the room.⁴⁴ Half of the chips were clearly labelled \$0, while the

⁴¹The experimental design called for 294 subjects - 98 per accuracy stake group. Sessions were run in all accuracy payment sessions until this minimum number (98) was reached. In one session (with 8 subjects) the experimental software crashed, leading to data for only one event. Including this session brings the total to 326.

⁴²Subject feedback suggested that they had a good grasp of the lottery method by the end of the practice section.

⁴³The quiz terminated automatically after five minutes, while the weather question prompted individuals for an answer at the end of two minutes. No subject took longer than two minutes to answer the weather question.

⁴⁴When the number of subjects was odd, an additional poker chip was placed in the bag to make the number even.

other half were labelled \$80. The probability of drawing $P = \$80$ was thus set at 50%, which corresponds to $\epsilon = 0.5$. One by one subjects would draw a chip from the bag, until every subject had a chip. The amount of money on the chip determined how much extra, P , they would earn if in the prize state and the event had occurred.⁴⁵

This random draw of chips was repeated before each of the four events, which made it clear that the drawing of P was independent across events. With the exception of the physical drawing of chips, the rest of the experiment took place on lab computers using the experimental software Z-Tree. Each chip had a unique code that would load the specified amount into the computer. After all subjects entered this code into the computer, the event was introduced to all subjects. Subjects were informed that the event did not change based on whether they drew a \$0 or an \$80 chip. They then proceeded to have their beliefs elicited about the event. This procedure was repeated four times, for each event.

After elicitation for all four events was complete, I came around one final time with two bags. The first bag contained an equal number of Red (meaning accuracy stage) and Blue (meaning prize stage) poker chips, with the total equal to the number of subjects.⁴⁶ The chip drawn from this bag determined whether a subject was paid for their decisions during the elicitation procedure (red, for which they could earn $a \in \{\$3, \$10, \$20\}$ depending on the session), or was paid according to the prize stage (blue) whereby they earned $\bar{a} = \$20$ automatically, plus the amount on their prize chip, $P \in \{\$0, \$80\}$. The second bag contained an equal number of chips for each of the four events. The number of chips in the bag was equal to the smallest multiple of four that was greater than or equal to the number of subjects.

Across the four events I intentionally chose to investigate belief bias in different domains, leading to differences across events. Of the four, the two dice events are the closest in similarity. Since in the model of Brunnermeier and Parker (2005) biased beliefs are constrained only by the accuracy payment x , one would expect that since both dice events have nearly identical objective probabilities, bias will be identical conditional on x and W . On the other hand, Bracha and Brown (2012) predict that the mental cost of distorting reality would be greater for the simpler dice event, hence conditional on x and W the bias on the simpler dice event should be smaller than on the more difficult dice event.

A direct comparison between the objective and more subjective events is less straightforward, and not possible without further restrictions on the models. Again Bracha and

⁴⁵Additionally that event had to have been randomly selected for payment, of the four.

⁴⁶Again when the number of subjects was odd, an additional chip was placed in the bag.

Brown (2012) predict a higher degree of bias for the subjective events, however this does not translate into a testable prediction for the experimental data.

For the two dice events, 50% of individuals were randomly assigned to have control over choosing their numbers, while the other half had their numbers chosen by a computer.⁴⁷ For the quiz event, 30% of subjects per session were randomly assigned to a control group that evaluated a random partner’s performance instead of their own.

5 Experimental Results

Basic demographic data was collected on subjects including gender, university major, age, as well as standardized questions from psychology designed to measure generalized optimism.⁴⁸ Summary statistic tables can be found in Appendix B.1, showing that covariates are balanced across both treatments.

I first present an overview of individual’s initial reports for the probability they believed a given event would occur, $\tilde{\pi}_i$. Table 1 presents summary statistics for individual priors for each of the four events, separating the quiz event into whether it involved one’s own score, or a random partner’s score. Here the 3x2 treatments are aggregated and summarized by event. The final row presents these summary reports for all events pooled together.

It is clear from Table 1 that there is fairly substantial bias across all events. For the two dice events, defining the “true” probability π is straightforward. For the weather event the actual proportion of subjects who estimated the correct temperature range, 25.77% was used. For the quiz $\pi = 15\%$ was used, which was the proportion of subjects who scored in the top 15% relative to a similar reference group.⁴⁹ Of course, this does not mean that every subject should report a prior of 15%, due to heterogeneity in test ability.

That all probability reports are biased upwards is consistent with the predictions of both models, and of optimism models in general. The qualitative predictions of Bracha and Brown (2012) appear to be borne out in this table. The average bias for the harder dice

⁴⁷Subjects either had control for both dice events, or no control for both events. The subjects were only aware of their own arrangement, and had no knowledge that any other arrangement existed.

⁴⁸After the experiment subjects were asked four questions taken from the Life Orientation Test - Revised (LOT-R) a revised version of a test used in psychology to distinguish generalized optimism versus pessimism. This revised version was developed and subsequently published by Scheier et al. (1994). Their original test involves 10 questions, however 4 are “fillers” which are not considered when constructing an index. Four of the six questions remaining were selected for the post-experiment questionnaire.

⁴⁹Subjects were informed the reference group consisted of students just like them, taking the same quiz.

Table 1

Summary Statistics About Reported Beliefs ($\tilde{\pi}_i$)

Event	“Truth” (π)	$\mu_{\tilde{\pi}}$	Upward Bias	$Q_{2,\tilde{\pi}}$	$\sigma_{\tilde{\pi}}$	N
Dice Easy	11.11	17.09	53.81%	11.00	15.47	318
Dice Hard	11.57	20.77	79.55%	15.00	18.04	318
Weather	25.77	63.23	145.36%	65.00	20.54	326
Quiz Self	15.00	50.18	234.53%	50.00	27.49	223
Quiz Other	15.00	26.69	77.96%	20.00	17.82	95
All	15.92	36.23	127.60%	30.00	27.91	1280

“Truth” refers to the baseline underlying probability of the event. $\mu_{\tilde{\pi}}$ is the mean of the belief report $\tilde{\pi}_i$, $Q_{2,\tilde{\pi}}$ is the median, and $\sigma_{\tilde{\pi}}$ is the standard deviation.

event is higher than that of the easier dice event, and the more subjective events appear to have larger degrees of bias. However, while these results are suggestive that individuals hold optimistic beliefs, they are by no means conclusive. Similar patterns could be generated by models of bounded rationality or costly effort, since errors are unlikely to be symmetrically distributed. In order to conclusively rule out such other explanations, I now turn to the results examining the comparative statics.

5.1 Comparative Static Predictions of Optimism Models

This section provides an overview of the empirical strategy, and results for testing the two major comparative static predictions of both BP and BB. For the empirical analysis, the dependent variable is the reported belief $\tilde{\pi}_{ij}$ of individual i for event j . $1 \leq j \leq 5$, an integer, is an indicator for the two dice events, the weather event, the quiz event (self) and finally the quiz event (other), respectively.

5.1.1 Testing Hypothesis 1: Does optimism bias increase with the prize stake P_{ij} ?

The first comparative static of interest is the coefficient β_1 on a dummy variable for whether the individual was given a prize stake $P_{ij} > 0$ for a given event j .

$$\tilde{\pi}_{ij} = \beta_1 \cdot 1\{P_{ij} > 0\} + \sum_{1 \leq j \leq 5} \gamma_j \cdot E_j + \eta \cdot S_i \cdot E_4 + \alpha_k + \epsilon_{ij} \quad (13)$$

The dependent variable $\tilde{\pi}_{ij}$ is the reported belief of the subject regarding the percent chance of an event occurring, coded between 0 and 100. The independent variable of interest is an indicator of whether or not the subject had an \$80 prize stake in the event, $1\{P_{ij} > 0\}$. E_j is an event specific fixed effect, α_k is a session level fixed effect, and ϵ_{ij} is an idiosyncratic error for each probability report. Standard errors are clustered at the individual level.

Heterogeneity in ability will be an important factor in the determination of the probability that one believes they scored in the top 15%. To account for this heterogeneity, I interact score on the quiz with the fixed effect for the quiz (self) event.⁵⁰ S_i is the subject's score on the quiz, $E_4 = 1$ only if the event was the quiz (self) event. Since the quiz was conducted before the experiment began, subjects did not know whether they would face $P = \$80$ or $P = \$0$, so there is no concern about endogeneity of P_{ij} .

Regarding the Weather event, there is no statistically significant relationship between how far off one's temperature estimate was from the truth and the probability that individual believed they were within 5 degrees of the truth so I do not include this interaction in the table.⁵¹

Recall that both models predict that individuals given a positive financial stake in an event will believe the event to be more likely to occur, i.e. $\beta_1 > 0$. This is in contrast to the standard RE model where $\beta_1 = 0$. I now examine whether there are patterns in the data consistent with the comparative statics in Hypothesis 1.

Table 2 examines the comparative static predictions for each of the three accuracy groups, as well as for the aggregate data. From this table it is evident that having a prize stake is only significant in the low accuracy payment ($a = \$3$) sessions, where β_1 is significant and positive at the 1% level. The effect is sizeable, and corresponds to a 5.82 percentage point increase in the prior probability reported by an individual, or a 17% increase from the mean probability report in those sessions. For the other accuracy

⁵⁰The fit is not substantially improved with non-linear estimates nor non-parametric estimates. Reassuringly the results are unchanged using a linear or non-parametric relationship so I control for test score using a linear relationship.

⁵¹Including it does not affect the results.

payment sessions, and the aggregate data, the effect of stakes is not significant. On the whole, there is evidence of optimism bias when individuals are given a prize stake in the event, but only when accuracy payments are low.

Looking at the coefficient on quiz score for the Quiz (self) event, overall scoring an additional point on the quiz resulted in an approximately 2.2 percentage point increase in the probability reported of scoring in the top 15%. Put another way, a one-standard deviation increase in test score results in a 5.8 percentage point increase in the belief about the probability of scoring in the top 15%. Interestingly, the relationship between quiz score and beliefs about scoring in the top 15% seems to be strongest at the lowest accuracy payment (\$3), and gets weaker as the accuracy payments increase. A similarly counterintuitive pattern of beliefs being less connected to reality can be observed looking at the event-level fixed effects: increasing accuracy payments is related to increases in average probability reports for these events. The next section examines Hypothesis 2 which will specifically address patterns of beliefs in response to variation in accuracy payment levels.

Table 2

Impact of Financial Prize on Beliefs (Testing Hypothesis 1)

All Events				
Regressor	Acc = \$3	Acc = \$10	Acc = \$20	All
$P = \$80 (\beta_1)$	5.824 (2.149)***	-0.190 (1.680)	-2.090 (2.158)	1.320 (1.190)
Easy Dice (γ_1)	10.636 (5.511)*	10.120 (3.757)***	16.665 (2.538)***	8.415 (3.659)**
Hard Dice (γ_2)	11.902 (5.776)**	12.310 (3.906)***	24.324 (2.667)***	12.075 (3.711)***
Weather (γ_3)	54.360 (5.944)***	55.737 (4.606)***	65.285 (2.820)***	54.438 (4.000)***
Quiz (Self) (γ_4)	24.040 (7.303)***	32.134 (5.877)***	60.096 (5.790)***	33.218 (4.806)***
Quiz (Other) (γ_5)	14.191 (6.100)**	15.776 (4.184)***	36.321 (4.496)***	17.982 (3.985)***
Score * Quiz (Self) (η)	3.685 (1.023)***	2.618 (0.959)***	-0.771 (1.251)	2.172 (0.647)***
Session Fixed Effects (α)	YES	YES	YES	YES
N	424	436	420	1280
R^2	0.490	0.556	0.529	0.513

Difference is significant at * 0.1; ** 0.05; *** 0.01. Robust Std Errors Clustered at Individual Level

Constant Omitted as $\sum_{j=1}^5 E_j = 1$. R^2 corrected for no-intercept.

5.1.2 Testing Hypothesis 2: Does optimism bias change with accuracy stakes?

I next turn to examining Hypothesis 2, regarding the accuracy payments a . Hypothesis 2 made opposing predictions for the two models regarding increases in accuracy payments. The BP model of Brunnermeier and Parker (2005) predicts that as the accuracy payment a increases, individuals will respond to these incentives and report less biased probabilities. The BB model of Bracha and Brown (2012) predicts that as a increases, individuals will actually report *more* biased probabilities as the likelihood of receiving the accuracy payment is increasing in the probability the event occurs.

The empirical estimation of Hypothesis 2 is similar to Equation (13), here β_2 is the coefficient of interest, demonstrating the relationship between accuracy incentive payment a and beliefs $\tilde{\pi}_{ij}$. Since a was randomized at the session level, I do not include session fixed effects.

$$\tilde{\pi}_{ij} = \beta_2 \cdot a + \sum_{1 \leq j \leq 5} \gamma_j \cdot E_j + \eta \cdot S_i \cdot E_4 + \epsilon_{ij} \quad (14)$$

Table 3 presents the results from the specification in Equation (14). Examining this table, one can see that increasing the accuracy payments tends to *increase* optimism, a prediction uniquely made by BB but one that is inconsistent with BP. Note that here it is possible that a fraction of subjects were reporting “pessimistic” priors ($\tilde{\pi}_{ij} < \pi_j$), such that optimism bias coincides with becoming more accurate.⁵² However, the comparative static predictions of both models are independent of the prior belief, such that Equation (14) is correctly specified for the question of interest. Nonetheless, in Appendix B.2 I examine the same question for only subjects with upwardly biased priors. Not surprisingly, given the sizeable bias in probability reports across all sessions (as seen in Table 1), optimism results in even further bias in probability reports.

Again it is important to control for Quiz score. Since subjects effort on the quiz could be affected by the accuracy payment, one concern is that these results are in part driven by higher test scores that resulted from higher effort exerted when accuracy payments were larger. However, the effect remains significant controlling for test scores, and in fact is still significant when the entire Quiz (Self) event is removed from the analysis.⁵³

⁵²In the data, only 22% of subjects had pessimistic or unbiased probability reports.

⁵³It is also reassuring that, as shown in Appendix B.1, test scores do not significantly differ across

The effect that increasing accuracy incentives leads to more biased probability reports seems to entirely be driven by events for which subjects had no prize stake. For this group, the average effect is quite substantial. Increasing the lottery method's payment by \$2.25 would result in a 1 percentage point increase in the reported belief. Increasing stakes from \$3 to \$20 has the effect of increasing beliefs by 7.5 percentage points. This is actually larger than the average effect of having a prize stake of \$80 in the event (which was 5.82 percentage points).

As a final note, the model in BP predicts that when there is No Stake, bias should not change as accuracy payments are varied. However this is rejected at the 1% level. Overall, the BP model cannot explain the patterns generated in the experiment, while the BB model comparative static predictions do match the observed data.

treatments.

Table 3

Impact of Accuracy Payment on Beliefs (Testing Hypothesis 2)

All Events			
Regressor	No Stake	Stake = \$80	All
Accuracy Payment (β_2)	0.444 (0.116)***	0.052 (0.145)	0.250 (0.094)***
Easy Dice (γ_1)	11.146 (1.575)***	17.660 (2.428)***	14.345 (1.491)***
Hard Dice (γ_2)	15.327 (1.811)***	20.810 (2.321)***	18.030 (1.503)***
Weather (γ_3)	57.772 (2.091)***	63.199 (2.234)***	60.493 (1.590)***
Quiz (Self) (γ_4)	35.878 (4.993)***	42.811 (4.416)***	39.400 (3.303)***
Quiz (Other) (γ_5)	21.429 (2.653)***	26.456 (2.951)***	23.967 (2.025)***
Score * Quiz (Self) (η)	2.367 (0.992)**	1.831 (0.838)**	2.110 (0.645)***
Session Fixed Effects	NO	NO	NO
N	646	634	1280
R^2	0.528	0.469	0.495

Difference is significant at * 0.1; ** 0.05; *** 0.01. Robust Std Errors Clustered at Individual Level

Constant Omitted as $\sum_{j=1}^5 E_j = 1$. R^2 corrected for no-intercept.

5.2 Discussion of Results

Looking at the comparative static predictions from Hypotheses 1 and 2, it is clear that the BB model seems to fit the pattern of results observed, while the BP model cannot account for these patterns. Providing subjects with larger payments for giving accurate responses does not reduce optimistic bias, but has the opposite effect. Such results provide evidence that the level of bias is not constrained by the financial costs that result from poor decisions taken in the experiment (i.e. the expected monetary loss from reporting an inaccurate belief). Instead the results support the theory of BB which emphasizes the direct costs of distorting reality in belief determination.

These results are consistent with the experiment of Mayraz (2013), who found evidence of optimistic belief formation when eliciting beliefs about hypothetical future stock prices. In his experiment, half of the sessions involved subjects whose income was increasing in future price, while the other half involved subjects whose income was decreasing in the price. Using a logarithmic scoring rule (LSR), he found that subjects were optimistically biased, and that this bias was larger when accuracy payments increased, as in the current study. Under the LSR, the BB model does not make unambiguous predictions for how optimistic bias will change, but to my knowledge it is the only model that is not inconsistent with the results of Mayraz (2013).

Looking back to the results in this paper, from the data in Tables 2 and 3, it is clear that not all treatments are positive and statistically significant, as would be predicted in the theory of BB. In particular, when accuracy payments are moderate (\$10) or high (\$20), from Table 2 one can see that the effect of giving a subject a $P = \$80$ stake in an event appears to have a slight negative effect, though not significantly different from zero. In Table 3, when the prize P is large (\$80) the estimated effect is positive, but very small and not statistically significant.

Small but positive findings for these treatments are to be expected if subjects are nearing the limit of their ability to distort reality. If this is true, because of the shape of the mental cost function $J^*(\cdot)$, increasing the benefits to optimistic belief formation will only result in minor changes to optimal beliefs. This can account for the findings in Table 3 (Column 2), but cannot explain the findings in Table 2 (Columns 2 and 3).

One explanation is that probability reports are subject to idiosyncratic error, or that there is unobserved heterogeneity in the mental cost function $J^*(\cdot)$.⁵⁴ A different expla-

⁵⁴In a between subjects design, if this heterogeneity is not balanced across treatments, it can lead to

nation can be found if one extends the model to allow subjects to also distort perceptions about the probability of ending up in the Accuracy State (ϵ) and distorting perceptions about the distribution used in the elicitation procedure. In Appendix A.2, I take steps to show how these comparative static results change when allowing agents more flexibility in what probabilities they can hold biased beliefs over. These results demonstrate how comparative static predictions of the BB model are unchanged when a is low and/or $P = \$0$. However it is possible that for $P = \$80$ and larger a , these results are no longer unambiguously positive, which may explain these findings in the data.

5.3 Alternative Explanations

In the previous section I argued that the results seem to be consistent with the model of Bracha and Brown (2012). This section is intended to examine whether the patterns observed might be consistent with other behavioral models.

First a number of explanations may be consistent with probability reports that are biased upwards. Because probabilities are naturally restricted to lie in $[0, 1]$, and because in the experimental design $\pi < \frac{1}{2}$, it is plausible that errors will be right skewed, consistent with the experimental results.

In particular, models of costly effort and/or attention might generate these patterns. However, they are unable to explain the comparative static results. Because costs of effort are structurally unrelated to accuracy payments in the experiment, while increasing payments for accuracy will increase the returns to effort, effort should increase, and probability reports should become more accurate. This is contrary to the observation that $\frac{d\hat{\pi}}{da} > 0$. Similarly for models of inattention, increasing accuracy payments a or prize payments P would be expected to increase attention, and thus decrease bias. This is the opposite of what is observed.

Two explanations that could potentially generate the comparative static findings are “Buckling under pressure” or “Satisficing”.

First, it is possible that individuals might make more errors when the stakes are higher, for example because they do not perform well under pressure. Thus, if subjects can earn \$20 from the lottery method, they feel pressure about earning this relatively large amount of money and this leads them to make mistakes. Such behavior could generate similar patterns to those observed in the experiment. However, I also have an independent measure

erroneous negative estimates (non-statistically significant), as observed in the data.

of performance under pressure, subjects' scores on the Quiz. And in fact, while there are no statistically significant differences between the three accuracy payment sessions, quiz scores are one half-point higher in the \$10 and \$20 sessions relative to the \$3 sessions. Such findings seem to rule out that subjects find it difficult to estimate probabilities under pressure.

The second explanation is that subjects might only have preferences over earning a minimum amount of money during the experiment, and once this reference level of income is (ex-ante) reached they may no longer find it worthwhile to exert effort. Thus for the \$20 accuracy sessions, subjects might have expectations of earning enough money such that they are content to exert little effort, while for the \$3 sessions they are faced with low expected earnings and are willing to exert more effort to increase their expected pay-out.

While I have no way to directly refute the predictions of such a model, it is possible to examine subject choice time. On average subjects in the \$3 and \$10 sessions took about 32 seconds to submit their probability report, while those in \$20 sessions took about 29 seconds, a non-statistically significant difference. Subjects did however spend about 5 seconds *longer* on their report when they had an \$80 Stake in the event, a difference significant at the 10% level. If choice time is a proxy for effort then it would seem that subjects did not exert less effort when the stakes or accuracy payments were higher, which such a model would suggest.

As a final comment on the viability of costly effort models to explain patterns in the data, I note that the experimental design did not provide any advantages to finishing any of the tasks faster. Subjects who made probability reports faster than their peers had to wait until the next task, and were not permitted to engage in any other activities. If there are costs of boredom, these should work to ensure that subjects were focussed on probability reports when they had the opportunity.

5.3.1 Other Models from Decision Theory

This experiment is also related to decision theoretic models that seek to explain deviations from expected utility. Focusing only on the Accuracy state, the lottery method used to elicit probability reports in this experiment is theoretically equivalent to pairwise comparisons of two lotteries. The first lottery pays out a if the event of interest happens, and 0 otherwise (denote this lottery by $(a, \hat{\pi})$). The second lottery is an objective lottery that pays out a with probability $r \in [0, 1]$, denoted by (a, r) . The subject must select an \underline{r} such that for

all values of $r \geq \underline{r}$ she would prefer (a, r) to $(a, \hat{\pi})$.

The behavioral result in this experiment is that \underline{r} changes with the value of a . In particular \underline{r} is increasing in a . This result is primarily driven in the condition when there was no prize stake, i.e. $P = 0$. For completeness the lottery is thus $(a, \epsilon \cdot \hat{\pi}; 20, 1 - \epsilon)$ versus $(a, \epsilon \cdot r; 20, 1 - \epsilon)$. Again the puzzle is that \underline{r} is increasing with a .

The interesting feature is that the loss function is symmetric, that is shifting \underline{r} higher results in the exact same expected loss as shifting \underline{r} lower by the same amount. This means that models of loss aversion or disappointment aversion would predict no change in decision making as a is varied.

On the other hand, certain classes of preferences that model ambiguity behavior can potentially explain these patterns. In particular, if subjects degree of ambiguity seeking is increasing in the financial stakes, then these preferences could generate the patterns observed in the experimental data. However, while there are models which are broad enough to be consistent with such preferences, such models are typically thought to apply to situations of ambiguity aversion or pessimism.⁵⁵ Instead, the model of BB is motivated precisely by the phenomenon of optimism, and has an axiomatization that makes it analogous to the variational preferences model of Maccheroni et al. (2006), with ambiguity-seeking rather than averse agents.

6 Conclusion

This experiment examined two models of optimistic belief formation, the optimal expectations model of Brunnermeier and Parker (2005) (BP), and the affective decision making model of Bracha and Brown (2012) (BB). Both models differ in how optimistic beliefs are constrained, highlighting intuitive mechanisms: BP states that biased beliefs lead to poor outcomes through sub-optimal decisions. A key feature of BB is that costs of distorting beliefs depends on the probabilistic environment.

The experiment exploits these differences and shows that optimistic beliefs are prevalent across different domains. When individuals are given a large financial stake in an event, I find some evidence that bias increases, consistent with both models. However, larger incentive payments for accurate reports leads to an *increase* in biased beliefs, a prediction made uniquely by BB. The BP model makes the opposite prediction, which is rejected in

⁵⁵An example is the model of Gumen et al. (2012).

the data.

Additionally these results are consistent with the related work of Mayraz (2013), who also found that higher accuracy payments led to a larger degree of bias.⁵⁶ The present experiment builds on this result, and shows that across a variety of events, with an incentive compatible elicitation procedure, this result is a more robust phenomenon, and can be explained in terms of the BB model.

The implications of these results are that direct costs of distorting beliefs matter for their determination. Oster et al. (2013) have shown that optimism is an important factor in understanding low rates of testing for Huntington Disease (HD). A question for policy is how to increase testing rates for the disease. This experiment has shown that there will be benefits to public health campaigns targeting at risk individuals, and educating them on their true risk status. In contrast, as Oster et al. (2013) discuss, the BP model provides only limited prescriptions for policy makers, such as drawing attention to how behavioral changes may lower the expected costs of having HD.

One question is whether the results of this experiment will hold more generally when the stakes are much higher, relative to those used in the lab. The costs of not being tested for HD likely substantially exceed the stakes in the experiment. Further work might show whether these results can be replicated for higher stakes.

Finally, these results shed important light on belief elicitation in experiments. Perhaps counter-intuitively, but good news for experimental economists, is that the conclusion of this experiment shows that beliefs will be most accurate precisely when incentives for accuracy are low. If incentives are too high, they may detract from other components of the experiment, and the incentive payments themselves may lead subjects to form optimistic beliefs.⁵⁷

⁵⁶In Mayraz (2013) the predictions of BB are not unambiguous as in the present experiment, however a positive association is not inconsistent with their model. The predictions of BP are rejected in Mayraz (2013) as well.

⁵⁷It is worth mentioning that there may be other reasons why higher incentive payments can be detrimental. Gächter and Renner (2010) show that incentivizing beliefs can alter behavior in public goods games.

Appendix

A Extension: Allowing Subjects to Distort Other Probabilities in the Model

In this section I relax the assumption that agents treat ϵ and the distribution used to determine the lottery method payout as given. Let $\hat{\epsilon}$ be the biased probability that an individual assigns to ϵ , and let \hat{F} be the distorted distribution that the random number r in the accuracy state is drawn from. The first step will be to show under what conditions truthful reporting is still optimal. In this respect I will make the following assumptions.

Assumption 1. *For a BP agent $\hat{\epsilon} \geq \underline{\epsilon} > 0$*

Assumption 2. *For a BP agent \hat{F} has full support on $[0, 1]$, and for all $r \in [0, 1]$, $\hat{F}(r) \geq \underline{F}(r)$, where \underline{F} is a well-behaved CDF with full support on $[0, 1]$ and $r \geq \underline{F}(r)$ for all $r \in [0, 1]$. Additionally the PDF of this distribution, \underline{f} is non-decreasing on $[0, 1]$.*

These assumptions are to ensure that a BP agent never believes that possible states of the world are impossible.⁵⁸ In order to guarantee a solution exists I assume there is some lower bound $\underline{\epsilon}$, which may in fact be extremely small, but non-zero. Similarly I assume that there exists a bound on how skewed the distribution \hat{F} can be, where \underline{F} is understood to be a highly right-skewed distribution that first order stochastically dominates the uniform distribution on $[0, 1]$. These assumptions will not be required for a BB agent since the mental cost function $J^*(\cdot)$ is such that the costs of holding such extreme beliefs will ensure such beliefs are never optimal.

Proposition 2. *Given Assumptions 1 and 2, an agent in the experiment truthfully reports her belief, regardless of whether she holds BP, BB, or RE beliefs.*

Proof. Here I refer to the proof of Proposition 1 for BB and RE agents. Consider only BP agents. At this stage the individual takes $\hat{\epsilon}$ and $\hat{\pi}$ as given, and maximizes first period utility. Assumption 1 guarantees that the report $\tilde{\pi}$ is relevant to the individual (if $\hat{\epsilon} = 0$ then the individual believes there is no chance of receiving payment for his report, and hence is at liberty to report anything). Assumption 2 guarantees that the lottery method

⁵⁸Intuitively the reason for Assumption 1 is that if an individual were 100% certain that her probability report in the experiment would not matter she no longer has any incentive to report truthfully. Similarly for Assumption 2 if she believed that there were gaps in the support of the distribution of drawing a random number r she would be indifferent between reporting any probability within such a gap.

is incentive compatible. With both these assumptions, the optimal report is $\tilde{\pi} = \hat{\pi}$ for a BP agent. \square

Given that an optimal expectations agent truthfully reports his subjective belief, I now turn to what optimal subjective beliefs $\hat{\pi}$ and $\hat{\epsilon}$ will be. This is separated into two sections, first for BP agents followed by BB agents.

A.1 Optimal Beliefs in Brunnermeier and Parker (2005)

First note that the prize state strictly dominates the accuracy state, since the prize state guarantees that wealth is at least \bar{a} regardless of whether E occurs or not, while in the accuracy state there is a probability of earning nothing, and the most the individual can earn is $a \leq \bar{a}$. Thus according to Brunnermeier and Parker (2005) the optimal biased belief for ϵ would be to select a belief as low as possible. Because of Assumption 1, the optimal belief will be $\hat{\epsilon}^{BP} = \underline{\epsilon}$.

Since $r \sim U[0, 1]$, the true probability that $r < \tilde{\pi}$ is $\tilde{\pi}$ and $E[r|r > \tilde{\pi}] = \frac{1+\tilde{\pi}}{2}$. For a BP agent believing that $r \sim \hat{F}[0, 1]$ the probability $r < \tilde{\pi} = \hat{F}(\tilde{\pi})$ and $E_{\hat{F}}[r|r > \tilde{\pi}] = \frac{\int_{\tilde{\pi}}^1 r d\hat{F}(r)}{1-\hat{F}(\tilde{\pi})}$.

Consider two distributions with full support on $[0, 1]$, \hat{F} and \hat{F}' .

Proposition 3. *If \hat{F} first order stochastically dominates \hat{F}' , the BP agent weakly prefers to believe \hat{F} rather than \hat{F}' .*

Proof. Given a BP agent truthfully reports his belief, the only relevance for this distribution is the expected value of r whenever $r > \hat{\pi}$. Thus the agent prefers that the probability $(1 - \hat{F}(\hat{\pi})) \cdot \frac{\int_{\hat{\pi}}^1 r d\hat{F}(r)}{1-\hat{F}(\hat{\pi})} = \int_{\hat{\pi}}^1 r d\hat{F}(r)$ be as large as possible. Since first order stochastic dominance ensures that $\int_{\hat{\pi}}^1 r d\hat{F}(r) \geq \int_{\hat{\pi}}^1 r d\hat{F}'(r)$ for any $\hat{\pi}$, the BP agent weakly prefers to believe \hat{F} . \square

Proposition 3 combined with Assumption 2 means that if the individual holds biased beliefs about the distribution of r , she will select the optimal biased belief that $r \sim \hat{F} = \underline{F}$.⁵⁹ Thus optimal beliefs will be $\hat{\epsilon}^{BP} = \underline{\epsilon}$ and $\hat{F}^{BP} = \underline{F}$, which critically, *do not* depend on the belief about the event E occurring, $\hat{\pi}$. We present the analog to Equations 8 and 9 in the more general framework. Utility at time $t = 1$ is:

⁵⁹ \underline{F} also has the property that \underline{f} is non-decreasing. Intuitively (and loosely speaking) a BP agent prefers such distributions.

$$\gamma \left[\underline{\epsilon} \cdot \left(\underline{F}(\tilde{\pi}) \cdot (\hat{\pi} \cdot u(a) + (1 - \hat{\pi}) \cdot u(0)) + (1 - \underline{F}(\tilde{\pi})) \cdot \left(\frac{\int_{\tilde{\pi}}^1 r d\underline{F}(r)}{1 - \underline{F}(\tilde{\pi})} \cdot u(a) + \left(1 - \frac{\int_{\tilde{\pi}}^1 r d\underline{F}(r)}{1 - \underline{F}(\tilde{\pi})}\right) \cdot u(0) \right) \right) \right. \\ \left. + (1 - \underline{\epsilon}) \cdot \left(\hat{\pi} \cdot u(P + \bar{a}) + (1 - \hat{\pi}) \cdot u(\bar{a}) \right) \right] \quad (15)$$

Utility at time $t = 2$ from second period consumption depends on the true probabilities and distributions, which is unchanged hence I reproduce Equation 9:

$$\epsilon \cdot \left(\tilde{\pi} \cdot (\pi \cdot u(a) + (1 - \pi) \cdot u(0)) + (1 - \tilde{\pi}) \cdot \left(\frac{1 + \tilde{\pi}}{2} \cdot u(a) + \left(1 - \frac{1 + \tilde{\pi}}{2}\right) \cdot u(0) \right) \right) + (1 - \epsilon) \cdot \left(\pi \cdot u(P + \bar{a}) + (1 - \pi) \cdot u(\bar{a}) \right)$$

As before I substitute $\tilde{\pi} = \hat{\pi}$, substituting this value in, the optimal choice of $\hat{\pi}$ is the maximization of Equations (15) + 9, and is implicitly solved by:

$$\gamma \underline{\epsilon} \cdot \underline{F}(\hat{\pi}^{BP}) - \hat{\pi}^{BP} \epsilon + \pi \epsilon + (1 - \underline{\epsilon}) \gamma \cdot \frac{u(P + \bar{a}) - u(\bar{a})}{u(a) - u(0)} = 0 \quad (16)$$

As in Equation 10 a solution to the maximization problem exists, however the solution may be at a corner. This includes the situation where $\gamma = 1$, $\underline{\epsilon} = \epsilon$ and $\underline{F}(\hat{\pi}^{BP}) = \hat{\pi}^{BP}$ which I ignore here since this situation was covered in the less general model. I now focus on finding an interior solution. Note that there may be up to two interior solutions, however the lower $\hat{\pi}^{BP}$ is the maximum.⁶⁰

Next I show that the comparative statics are identical to the less general model. First consider the derivative with respect to W :

⁶⁰The second order condition is:

$$\gamma \underline{\epsilon} \cdot \underline{f}(\hat{\pi}^{BP}) - \epsilon \quad (17)$$

$\underline{f}(0) = 0$ and by Assumption 2 \underline{f} is non-decreasing. Hence there is one inflection point, there are at most two interior optima, the first of which is a maximum while the second is a minimum. If there is only one interior optimum it will be a maximum.

$$\frac{d\hat{\pi}^{BP}}{dP} = \frac{(1 - \underline{\epsilon})\gamma \cdot \frac{u'(P+\bar{a})}{u(a)-u(0)}}{\epsilon - \gamma\underline{\epsilon} \cdot \underline{f}(\hat{\pi}^{BP})} > 0 \quad (18)$$

Next consider the derivative with respect to x :

$$\frac{d\hat{\pi}^{BP}}{da} = -\frac{(1 - \underline{\epsilon})\gamma \cdot \frac{u(P+\bar{a})-u(\bar{a})}{(u(a)-u(0))^2} u'(a)}{\epsilon - \gamma\underline{\epsilon} \cdot \underline{f}(\hat{\pi}^{BP})} < 0 \quad (19)$$

Where the denominator in both expressions is positive from the second order conditions for a maximum. Thus, even when allowing a BP agent to distort the probability about being in the accuracy or prize state (ϵ) or to distort the distribution from which the lottery method draws a random number from $r \sim F$ the comparative static results continue to hold.

A.2 Optimal Beliefs in Bracha and Brown (2012)

First, I make an assumption regarding beliefs in Bracha and Brown (2012) regarding separability of the mental costs of distorting beliefs. Namely I assume that:

Assumption 3. *If two events E_1 and E_2 are stochastically independent, then the mental cost of distorting the probability of an outcome in E_1 , $J_{E_1}^*(\cdot)$ is independent of the mental cost of distorting the probability of an outcome in E_2 , $J_{E_2}^*(\cdot)$. Moreover, I require that changes in the event E_1 do not affect the cost function $J_{E_2}^*(\cdot)$, and vice-versa, so that the mental cost functions are additively separable in the objective function of the emotional process.*

Recall again that $J^*(\hat{x}, x)$ is strictly convex, which will ensure uniqueness. I now turn to how a Bracha and Brown (2012) agent might behave if she also distorts her perception of the distribution under which the lottery method operates, as well as the probability of the accuracy state ϵ . In order to allow this I must make some restrictions in order to make the theory tractable. In particular I assume that agents believe that r used in the lottery method is drawn from a beta distribution, with parameters $\hat{\alpha}$ and $\beta = 1$. In this way, $\hat{\alpha} = \alpha = 1$ corresponds to the uniform distribution, which is the “true” distribution used. I assume the standard Bracha and Brown (2012) cost function, but a slightly different

interpretation. α is a parameter, not a probability, but this interpretation allows a reduced form simplification that is tractable.

Specifically $J_\alpha^*(\hat{\alpha}, \alpha)$ takes a minimum at $\hat{\alpha} = \alpha = 1$ and $\lim_{\hat{\alpha} \rightarrow 0} J_\alpha^*(\hat{\alpha}, \alpha) = \lim_{\hat{\alpha} \rightarrow K} J_\alpha^*(\hat{\alpha}, \alpha) = \infty$ for some $K > 1$. The only change in defining this function is that previously $K = 1$ whereas now $K > 1$. The beta distribution is particularly intuitive for agents with biased beliefs as the PDF of the beta distribution with $\beta = 1$ is monotonic. When good outcomes are paid out from a lottery with probability $r \sim \text{Beta}(\hat{\alpha}, 1)$, optimistic agents prefer $\hat{\alpha} > 1$, while pessimistic agents would prefer $\hat{\alpha} < 1$.

The updated objective function of the emotional process (U) is (where the PDF of $\text{Beta}(\hat{\alpha}, 1)$ is $\hat{\alpha}x^{\hat{\alpha}}$, and the CDF is $x^{\hat{\alpha}}$):

$$\begin{aligned} \max_{\hat{\pi}, \hat{\epsilon}, \hat{\alpha}} U(\hat{\pi}, \hat{\epsilon}, \hat{\alpha}, P, a, \bar{a}) = \max_{\hat{\pi}, \hat{\epsilon}, \hat{\alpha}} & \left\{ \hat{\epsilon} \cdot \tilde{\pi}^{\hat{\alpha}} \cdot \left(\hat{\pi} \cdot u(a) + (1 - \hat{\pi}) \cdot u(0) \right) + \right. \\ & \left. \left(1 - \tilde{\pi}^{\hat{\alpha}} \right) \cdot \left(\frac{\hat{\alpha}}{\hat{\alpha} + 1} \cdot \frac{(1 - \tilde{\pi}^{\hat{\alpha} + 1})}{1 - \tilde{\pi}^{\hat{\alpha}}} \cdot u(a) + \left(1 - \frac{\hat{\alpha}}{\hat{\alpha} + 1} \cdot \frac{(1 - \tilde{\pi}^{\hat{\alpha} + 1})}{1 - \tilde{\pi}^{\hat{\alpha}}} \right) \cdot u(0) \right) \right] + \\ & \left. (1 - \hat{\epsilon}) \cdot \left[\hat{\pi} \cdot u(P + \bar{a}) + (1 - \hat{\pi}) \cdot u(\bar{a}) \right] - J_\pi^*(\hat{\pi}, \pi) - J_\epsilon^*(\hat{\epsilon}, \epsilon) - J_\alpha^*(\hat{\alpha}, \alpha) \right\} \end{aligned}$$

Next I solve for the three FOCs and substitute in $\hat{\pi}$ for $\tilde{\pi}$. Thus these three first order conditions hold with equality, and make up the equilibrium of the BB model.

First, the FOC for $\hat{\pi}$ is:

$$\hat{\epsilon} \cdot \hat{\pi}^{\hat{\alpha}} \cdot \left(u(a) - u(0) \right) + (1 - \hat{\epsilon}) \cdot \left(u(P + \bar{a}) - u(\bar{a}) \right) - J_\pi^*(\hat{\pi}, \pi) = 0$$

The FOC for $\hat{\epsilon}$ is:

$$\frac{\hat{\alpha} + \hat{\pi}^{\hat{\alpha} + 1}}{1 + \hat{\alpha}} \cdot \left(u(a) - u(0) \right) + u(0) - \left[\hat{\pi} \cdot u(P + \bar{a}) + (1 - \hat{\pi}) \cdot u(\bar{a}) \right] - J_\epsilon^*(\hat{\epsilon}, \epsilon) = 0$$

Finally the FOC for $\hat{\alpha}$ is:

$$\hat{\epsilon} \cdot \left[\left(\frac{1 - \hat{\pi}^{\hat{\alpha} + 1}}{(1 + \hat{\alpha})^2} + \frac{\log \hat{\pi} \cdot \hat{\pi}^{\hat{\alpha} + 1}}{1 + \hat{\alpha}} \right) \cdot \left(u(a) - u(0) \right) \right] - J_\alpha^*(\hat{\alpha}, \alpha) = 0$$

Next, it will be useful for notation to use the following cross partial derivatives where possible for the analysis:

$$\begin{aligned}\frac{d^2U}{d\hat{\pi}d\hat{\epsilon}} &= \hat{\pi}^{\hat{\alpha}} \cdot \left(u(a) - u(0) \right) - \left(u(P + \bar{a}) - u(\bar{a}) \right) \\ \frac{d^2U}{d\hat{\pi}d\hat{\alpha}} &= \hat{\epsilon} \cdot \log \hat{\pi} \cdot \hat{\pi}^{\hat{\alpha}} \cdot \left(u(a) - u(0) \right) < 0 \\ \frac{d^2U}{d\hat{\epsilon}d\hat{\alpha}} &= \left(\frac{1 - \hat{\pi}^{\hat{\alpha}+1}}{(1 + \hat{\alpha})^2} + \frac{\log \hat{\pi} \cdot \hat{\pi}^{\hat{\alpha}+1}}{1 + \hat{\alpha}} \right) \cdot \left(u(a) - u(0) \right) > 0\end{aligned}$$

I next totally differentiate all three equations with respect to the accuracy payment a . For notation, I denote G_π as the equation from the FOC with respect to $\hat{\pi}$, and similarly G_ϵ and G_α for the other two. This differentiation gives:

$$\begin{aligned}\frac{dG_\pi}{da} &= \hat{\epsilon} \cdot \hat{\pi}^{\hat{\alpha}} \cdot u'(a) + \frac{d^2U}{d\hat{\pi}d\hat{\alpha}} \cdot \frac{d\hat{\alpha}}{da} + \frac{d^2U}{d\hat{\pi}d\hat{\epsilon}} \cdot \frac{d\hat{\epsilon}}{da} + \frac{d^2U}{d^2\hat{\pi}} \cdot \frac{d\hat{\pi}}{da} = 0 \\ \frac{dG_\epsilon}{da} &= u'(a) \cdot \left(\frac{\hat{\alpha}}{\hat{\alpha} + 1} + \frac{1}{\hat{\alpha} + 1} \hat{\pi}^{\hat{\alpha}+1} \right) + \frac{d^2U}{d\hat{\epsilon}d\hat{\alpha}} \cdot \frac{d\hat{\alpha}}{da} + \frac{d^2U}{d\hat{\pi}d\hat{\epsilon}} \cdot \frac{d\hat{\pi}}{da} - J_\epsilon^{*''} \cdot \frac{d\hat{\epsilon}}{da} = 0 \\ \frac{dG_\alpha}{da} &= \hat{\epsilon} \cdot \left(\frac{1 - \hat{\pi}^{\hat{\alpha}+1}}{(1 + \hat{\alpha})^2} + \frac{\log \hat{\pi} \cdot \hat{\pi}^{\hat{\alpha}+1}}{1 + \hat{\alpha}} \right) \cdot u'(a) + \frac{d^2U}{d\hat{\epsilon}d\hat{\alpha}} \cdot \frac{d\hat{\epsilon}}{da} + \frac{d^2U}{d\hat{\pi}d\hat{\alpha}} \cdot \frac{d\hat{\pi}}{da} + \frac{d^2U}{d^2\hat{\alpha}} \cdot \frac{d\hat{\alpha}}{da} = 0\end{aligned}$$

The solution to this system of equations gives the comparative static for $\frac{d\hat{\pi}^{BB}}{da}$:

$$\begin{aligned}\frac{d\hat{\pi}^{BB}}{da} &= \mathcal{D}_\pi^{-1} \cdot \left[\hat{\epsilon} \cdot \hat{\pi}^{\hat{\alpha}} \cdot u'(x) + \frac{d^2U}{d\hat{\pi}d\hat{\alpha}} \cdot \frac{\hat{\epsilon} \cdot \left(\frac{1 - \hat{\pi}^{\hat{\alpha}+1}}{(1 + \hat{\alpha})^2} + \frac{\log \hat{\pi} \cdot \hat{\pi}^{\hat{\alpha}+1}}{1 + \hat{\alpha}} \right) \cdot u'(x)}{-\frac{d^2U}{d^2\hat{\alpha}}} + \right. \\ &\quad \left. \left(\frac{\frac{d^2U}{d\hat{\pi}d\hat{\alpha}} \cdot \frac{d^2U}{d\hat{\epsilon}d\hat{\alpha}}}{-\frac{d^2U}{d^2\hat{\alpha}}} + \frac{d^2U}{d\hat{\pi}d\hat{\epsilon}} \right) \cdot \frac{u'(x) \cdot \hat{\pi}^{\hat{\alpha}+1} + \frac{d^2U}{d\hat{\epsilon}d\hat{\alpha}} \cdot \frac{\hat{\epsilon} \cdot \left(\frac{1 - \hat{\pi}^{\hat{\alpha}+1}}{(1 + \hat{\alpha})^2} + \frac{\log \hat{\pi} \cdot \hat{\pi}^{\hat{\alpha}+1}}{1 + \hat{\alpha}} \right) \cdot u'(x)}{-\frac{d^2U}{d^2\hat{\alpha}}}}{J_\epsilon^{*''} - \frac{d^2U}{d\hat{\epsilon}d\hat{\alpha}} \cdot \left(-\frac{d^2U}{d^2\hat{\alpha}} \right)^{-1}} \right]\end{aligned}$$

$$\text{Where } \mathcal{D}_\pi = -\frac{d^2U}{d^2\hat{\pi}} - \frac{d^2U}{d\hat{\pi}d\hat{\alpha}} \cdot \frac{d^2U}{d\hat{\epsilon}d\hat{\alpha}} - \left(\frac{\frac{d^2U}{d\hat{\pi}d\hat{\alpha}} \cdot \frac{d^2U}{d\hat{\epsilon}d\hat{\alpha}}}{-\frac{d^2U}{d^2\hat{\alpha}}} + \frac{d^2U}{d\hat{\pi}d\hat{\epsilon}} \right) \cdot \left(J_\epsilon^{*''} - \frac{d^2U}{d\hat{\epsilon}d\hat{\alpha}} \cdot \left(-\frac{d^2U}{d^2\hat{\alpha}} \right)^{-1} \right)^{-1} > 0.$$

The sign of the denominator follows from the second order conditions.

$$\text{Letting } \kappa = \frac{\frac{1-\hat{\pi}^{\hat{\alpha}+1}}{(1+\hat{\alpha})^2} + \frac{\log \hat{\pi} \cdot \hat{\pi}^{\hat{\alpha}+1}}{1+\hat{\alpha}}}{-\frac{d^2U}{d^2\hat{\alpha}}},$$

$$\frac{d\hat{\pi}^{BB}}{da} = \left(\mathcal{D}_\pi \cdot u'(x) \right)^{-1} \left[\hat{\epsilon} \cdot \hat{\pi}^{\hat{\alpha}} + \frac{d^2U}{d\hat{\pi}d\hat{\alpha}} \cdot \hat{\epsilon} \cdot \kappa + \frac{\frac{d^2U}{d\hat{\pi}d\hat{\alpha}} \cdot \frac{d^2U}{d\hat{\epsilon}d\hat{\alpha}} + \frac{d^2U}{d\hat{\pi}d\hat{\epsilon}}}{-\frac{d^2U}{d^2\hat{\alpha}}} \cdot \left(\hat{\pi}^{\hat{\alpha}+1} + \frac{d^2U}{d\hat{\epsilon}d\hat{\alpha}} \cdot \hat{\epsilon} \cdot \kappa \right) \right] \\ J_\epsilon^{*''} - \frac{d^2U}{d\hat{\epsilon}d\hat{\alpha}} \cdot \left(-\frac{d^2U}{d^2\hat{\alpha}} \right)^{-1}$$

While this term is no longer unambiguously positive as in the main paper, I will make an assumption that will restore this earlier result when $P = 0$. I assume:

Assumption 4. $\hat{\pi} \geq \underline{\pi} = 0.016$ and $J_\alpha^{*''} > \hat{\epsilon} \cdot (u(a) - u(0))$.

The first part of this assumption requires that the biased probability $\hat{\pi}$ be not *too* small. 1.6% is an upper bound on the restriction, assuming that all other choice probabilities $\hat{\epsilon}$ and $\hat{\alpha}$ are stacked against giving the comparative static result. In the context of the experiment, this assumption is not restrictive at all. Only 24/1280 observations report a belief less than 1.6%. $\hat{\pi} \geq \underline{\pi}$ is required to help ensure that the third term in the square brackets is positive, when $P = 0$. The second part of the assumption requires that the second derivative of the mental cost function for distorting α is as large as $\hat{\epsilon} \cdot (u(a) - u(0))$. This assumption is more difficult to interpret, but provides a relatively simple way of ensuring that the first term in the comparative static is larger than the absolute value of the second negative term, and guarantees that the third term is positive. Clearly, this condition is more likely to be satisfied when a is smaller.

Given Assumption 4, the comparative static result $\frac{d\hat{\pi}}{da} > 0$ continues to hold, even permitting agents to distort the probability of ϵ and the distribution used in the elicitation procedure. Note again, that Assumption 4 is a sufficient but not necessary condition for this result, and only guarantees this result when when $P = 0$.

Finally, I am similarly able to compute the second comparative static $\frac{d\hat{\pi}^{BB}}{dP}$.

$$\frac{d\hat{\pi}^{BB}}{dP} = \left(\mathcal{D}_\pi \right)^{-1} \left[(1 - \hat{\epsilon}) \cdot u'(P + \bar{a}) - \hat{\pi} \cdot u'(P + \bar{a}) \cdot \frac{\frac{d^2U}{d\hat{\pi}d\hat{\epsilon}} + \frac{d^2U}{d\hat{\pi}d\hat{\alpha}} \cdot \frac{d^2U}{d\hat{\epsilon}d\hat{\alpha}}}{-\frac{d^2U}{d^2\hat{\alpha}}} \right] \\ J_\epsilon^{*''} - \frac{d^2U}{d\hat{\epsilon}d\hat{\alpha}} \cdot \left(-\frac{d^2U}{d^2\hat{\alpha}} \right)^{-1}$$

Here, when $P = 80$, it is likely that $\frac{d^2U}{d\hat{\pi}d\hat{e}} < 0$ ⁶¹. In this case the final term multiplying $-\hat{\pi} \cdot u'(P + \bar{a})$ will be negative, making the entire second term positive. As a result $\frac{d\hat{\pi}^{BB}}{dP} > 0$. In contrast, when $P = 0$, given Assumption 4, the final term multiplying $-\hat{\pi} \cdot u'(P + \bar{a})$ will be positive. However, this term can also be shown to be less than $\hat{\pi}$ in absolute value. Thus one further condition is sufficient to always ensure that this comparative static will be positive, namely:

$$1 - \hat{e} \geq \hat{\pi}^2$$

Since $\hat{e} \leq \frac{1}{2}$, since individuals prefer the prize state, one can guarantee the result for $\hat{\pi} \leq 0.707$. However, it is important to note that the size of the term multiplying $-\hat{\pi} \cdot u'(P + \bar{a})$ is *decreasing* in a , such that the comparative static is extremely likely to hold for low values of a , but is slightly less likely to hold when a is large. This could potentially explain the negative coefficients of the estimated data.

To summarize the theory in this section, when extending the models to allow individuals to distort other probabilities, the results continue to go through under some assumptions. For the BP model, the key assumption is that individuals must not believe that states of the world are impossible. For BB the key assumption required is that $J_{\alpha}^{*''} > \hat{e} \cdot (u(a) - u(0))$. This assumption is most likely to hold when a is low, and perhaps less likely to hold when a is larger. This will guarantee the comparative static predictions for changing a when $P = 0$, the treatment that saw the largest effect in the data. For the comparative static for P , it is more likely to hold for low a as well. These theoretical predictions match the findings in the data, and offer a different explanation of why some estimated coefficients are negative in the results.

B Supplemental Tables

B.1 Summary Statistics

Table 4 describes summary statistics at the *event* level, comparing the No Stake versus \$80 Stake condition, in order to check that the covariates are balanced across both treatments. The table shows that there are no significant differences across any of the observed variables. The table also details information about participants, approximately 41% of subjects are

⁶¹If it is not negative, the rest of the paragraph explains the conditions necessary for it to be positive.

male and 19% are economics or mathematics majors.

Table 4

Summary Statistics: Randomization of $P \in \{\$0, \$80\}$

	\$0 Stake	\$80 Stake	Diff	Total
Male	0.399	0.415	-0.016	0.407
Age	20.788	20.639	0.148	20.714
Test Score	3.931	3.704	0.227	3.821
Econ/Math Major	0.180	0.203	-0.023	0.192
Optimism Index	13.082	13.054	0.028	13.068
\$100 gamble	0.696	0.713	-0.017	0.705
N	642	630	-	1272†

Difference is significant at * 0.1; ** 0.05; *** 0.01 (Std. Errors clustered at Individual level.)

† N varies by demographic variable, e.g. as low as N=1220 for Male.

Variable Descriptions: Test Score is absolute score on quiz (max score 10, minimum -5. Optimism index from 0-16, higher indicates more optimistic on LOT-R style test. \$100 Gamble indicates preference for 50-50 chance at \$100 over \$20 with certainty.

Table 5 examines these same summary statistics grouped according to the accuracy payments at the session level. Here the observation unit is the individual which accounts for why the number of observations differs from the previous table. Note that the two variables “Test Score” and “Optimism Index” could vary endogenously with the level of accuracy incentives, as subjects may alter their effort on the test or change their reported optimism level depending on the size of the accuracy payment. Nonetheless, there appears to be no differences in any of these variables across the three accuracy payment treatments.

The last four columns examine whether there are significant differences between any of the three groups, while the final column test for equality of means across all three. None of the differences are significant at conventional levels, but note that the gender composition is close to being significantly different at the 10% level between the \$3 and \$10 groups, as is the Optimism Index between the \$10 and \$20 groups. Overall the randomization at the session level appears to have been successful, as the observed covariates are approximately

Table 5

Summary Statistics: Randomization of Accuracy Payments

<i>Accuracy:</i>	<i>\$3</i>	<i>\$10</i>	<i>\$20</i>	All	Diff (3,10)	Diff (10,20)	Diff (3,20)	P-Value All
Male	0.467	0.354	0.394	0.407	0.112	-0.040	0.072	0.258
Age	20.858	20.753	20.533	20.714	0.106	0.219	0.325	0.346
Test Score	3.538	3.966	3.962	3.821	-0.429	0.004	-0.424	0.402
Econ/Math Major	0.198	0.206	0.171	0.192	-0.008	0.035	0.027	0.799
Optimism Index	13.217	12.639	13.314	13.068	0.578	-0.675	-0.097	0.269
\$100 gamble	0.689	0.680	0.743	0.705	0.008	-0.062	-0.054	0.556
N†	106	107	105	318	-	-	-	-

Difference is significant at * 0.1; ** 0.05; *** 0.01. P-Value for multiple sample test of means (allows heterogeneous covariance).

† N varies slightly by demographic variable.

Variable Descriptions: Test Score is absolute score on quiz (max score 10, minimum -5. Optimism index from 0-16, higher indicates more optimistic on LOT-R style test. \$100 Gamble indicates preference for 50-50 chance at \$100 over \$20 with certainty.

balanced across different sessions.

B.2 Hypothesis 2: Biased Priors Only

Here Table 6 examines Hypothesis 2 restricting the sample to only individuals who reported upwardly biased probability reports. This examines the question of whether the optimistic bias observed in Table 3 is partly accounted for by subjects becoming more accurate. In fact, it can be seen that the results are very similar only looking at the biased sample, indicating that the optimism bias is leading subjects to report probabilities even further from the truth as accuracy payments increase.

Table 6

Impact of Accuracy Payment on Beliefs (Testing Hypothesis 2)

All Events: Only individuals who reported biased priors			
Regressor	No Stake	Stake = \$80	All
Accuracy Payment (β_2)	0.474 (0.122)***	-0.038 (0.162)	0.215 (0.107)**
Easy Dice (γ_2)	20.915 (1.965)***	33.238 (3.418)***	26.753 (2.001)***
Hard Dice (γ_2)	26.021 (2.276)***	31.216 (2.825)***	28.622 (1.886)***
Weather (γ_3)	60.476 (2.001)***	66.925 (2.337)***	63.792 (1.644)***
Quiz (Self) (γ_4)	43.116 (4.804)***	49.946 (4.545)***	46.529 (3.320)***
Quiz (Other) (γ_5)	23.387 (2.749)***	31.301 (3.268)***	27.558 (2.223)***
Score * Quiz (Self) (η)	1.427 (0.966)	1.084 (0.815)	1.288 (0.633)**
Session Fixed Effects	NO	NO	NO
N	464	459	923
R^2	0.472	0.394	0.425

Difference is significant at * 0.1; ** 0.05; *** 0.01. Robust Std Errors Clustered at Individual Level

Constant Omitted as $\sum_{j=1}^5 E_j = 1$. R^2 corrected for no-intercept.

Sample restricted to include only individuals who reported optimistically biased priors.

References

- Akerlof, George A and William T Dickens**, “The Economic Consequences of Cognitive Dissonance,” *American Economic Review*, February 1982, *72* (3), 307–319.
- Armantier, Olivier and Nicolas Treich**, “Eliciting beliefs: Proper scoring rules, incentives, stakes and hedging,” *European Economic Review*, August 2013, *62*, 17–40.
- Benabou, R. and J. Tirole**, “Self-Confidence and Personal Motivation,” *The Quarterly Journal of Economics*, August 2002, *117* (3), 871–915.
- Blanco, Mariana, Dirk Engelmann, Alexander K. Koch, and Hans-Theo Normann**, “Belief elicitation in experiments: is there a hedging problem?,” *Experimental Economics*, July 2010, *13* (4), 412–438.
- Bracha, Anat and Donald J. Brown**, “Affective decision making: A theory of optimism bias,” *Games and Economic Behavior*, May 2012, *75* (1), 67–80.
- Bridet, Luc and Peter Schwardmann**, “Selling dreams : Endogenous overoptimism and collateral use in financial contracts,” *mimeo*, 2014, pp. 1–62.
- Brunnermeier, Markus K and Jonathan A Parker**, “Optimal Expectations,” *American Economic Review*, 2005, *95* (4), 1092–1118.
- Caplin, Andrew and John Leahy**, “Psychological Expected Utility Theory and Anticipatory Feelings,” *Quarterly Journal of Economics*, 2001, *116* (1), 55–79.
- Easterwood, John C and Stacey R Nutt**, “Inefficiency in Analysts’ Earnings Forecasts: Systematic Misreaction or Systematic Optimism?,” *The Journal of Finance*, October 1999, *54* (5), 1777–1797.
- Eil, David and Justin M Rao**, “The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself,” *American Economic Journal: Microeconomics*, May 2011, *3* (2), 114–138.
- Ertac, Seda**, “Does self-relevance affect information processing? Experimental evidence on the response to performance and non-performance feedback,” *Journal of Economic Behavior & Organization*, December 2011, *80* (3), 532–545.

- Etner, Johanna, Meglena Jeleva, and Jean-Marc Tallon**, “Decision Theory Under Ambiguity,” *Journal of Economic Surveys*, April 2012, *26* (2), 234–270.
- Fischbacher, Urs**, “z-Tree: Zurich toolbox for ready-made economic experiments,” *Experimental Economics*, February 2007, *10* (2), 171–178.
- Gächter, Simon and Elke Renner**, “The effects of (incentivized) belief elicitation in public goods experiments,” *Experimental Economics*, July 2010, *13* (3), 364–377.
- Gilboa, Itzhak and David Schmeidler**, *A Theory of Case-Based Decisions*, Cambridge: Cambridge University Press, 2001.
- Grossman, Zachary and David Owens**, “An unlucky feeling: Overconfidence and noisy feedback,” *Journal of Economic Behavior & Organization*, November 2012, *84* (2), 510–524.
- Gumen, Anna, Efe Ok, and Andrei Savochkin**, “Decision-Making under Subjective Risk : Toward a General Theory of Pessimism,” 2012.
- Karni, Edi**, “A Mechanism for Eliciting Probabilities,” *Econometrica*, 2009, *77* (2), 603–606.
- **and Zvi Safra**, “The impossibility of experimental elicitation of subjective probabilities,” *Theory and Decision*, May 1995, *38* (3), 313–320.
- Koszegi, Botond**, “Health anxiety and patient behavior.,” *Journal of health economics*, November 2003, *22* (6), 1073–84.
- Köszegi, Botond**, “Ego Utility, Overconfidence, and Task Choice,” *Journal of the European Economic Association*, June 2006, *4* (4), 673–707.
- Koszegi, Botond**, “Utility from anticipation and personal equilibrium,” *Economic Theory*, May 2010, *44* (3), 415–444.
- Landier, Augustin**, “Wishful thinking a model of optimal reality denial,” 2000.
- **and David Thesmar**, “Financial Contracting with Optimistic Entrepreneurs,” *Review of Financial Studies*, April 2009, *22* (1), 117–150.

- Langer, Ellen J**, “The illusion of control.,” *Journal of Personality and Social Psychology*, 1975, *32* (2), 311–328.
- Loewenstein, George**, “Anticipation and the valuation of delayed consumption,” *The Economic Journal*, 1987, *97* (387), 666–684.
- Maccheroni, Fabio, Massimo Marinacci, and Aldo Rustichini**, “Ambiguity Aversion, Robustness, and the Variational Representation of Preferences,” *Econometrica*, November 2006, *74* (6), 1447–1498.
- Mayraz, Guy**, “Wishful Thinking,” *mimeo*, 2013, pp. 1–34.
- Mobius, Markus, Muriel Niederle, Tanya S Rosenblat, and Paul Niehaus**, “Managing Self-Confidence : Theory and Experimental Evidence,” *mimeo*, 2013.
- Oster, Emily, Ira Shoulson, and E. Ray Dorsey**, “Optimal Expectations and Limited Medical Testing: Evidence from Huntington Disease,” *American Economic Review*, April 2013, *103* (2), 804–830.
- Puri, M and D Robinson**, “Optimism and economic choice,” *Journal of Financial Economics*, October 2007, *86* (1), 71–99.
- Scheier, Michael F, Charles S Carver, and Michael W Bridges**, “Distinguishing Optimism From Neuroticism (and Trait Anxiety , Self-Mastery , and Self-Esteem): A Reevaluation of the Life Orientation Test,” 1994, *67* (6), 1063–1078.
- Spiegler, Ran**, “On Two Points of View Regarding Revealed Preferences and Behavioral Economics,” 2008, pp. 95–115.
- Vosgerau, Joachim**, “How prevalent is wishful thinking? Misattribution of arousal causes optimism and pessimism in subjective probabilities.,” *Journal of experimental psychology. General*, February 2010, *139* (1), 32–48.