Rational self-deception

Eli Spiegelman [a]

Jan 31, 2018

[a] CEREN, EA 7477, Burgundy School of Business - Université Bourgogne Franche-Comté, 29 rue Sambin, Dijon 21000, France. eli.spiegelman@bsb-education.com

Keywords: Wishful thinking; self-deception; beliefs; choice theory; experiment; regret theory; ambiguity

JEL D8, D9

Funding This research benefitted from funding from LESSAC, Burgundy School of Business.

Abstract

This note presents a model of wishful thinking as manipulation of prior beliefs, and a novel experimental test. The theoretical context is a three-period model, in which the agent is faced with a state-contingent optimal action, in which one state yields a higher payoff. In period 0 she observes the objective prior probability that each state will occur, but may alter her beliefs about these probabilities (self-deceive). The beliefs she chooses in period 0 determine her action in period 1 as a standard maximization procedure. In period 2, a signal yields information about the state of the world. A key assumption is that this signal may not be perfectly revealing. It is shown that the objective prior is optimal if and only if the signal in period 2 is perfectly revealing. Predictions of the theory are tested in a bet-choice experiment. I then present an experimental test designed to investigate the model's predictions. Subjects choose to play a bet or pass it up for another, and the experimental control varies whether learn about the second bet upon taking the first. The results of the experiment are also compared to the literature on regret theory and ambiguity aversion.

1. Introduction

Jack recently saw a movie about the production of industrial chicken. The movie showed, in graphic and (possibly intentionally) horrifying detail, the process by which agro-industrial firms maximize the flow of meat. Jack is recalling the movie now, in the grocery store, while deliberating between a conventional chicken ($9.99) and an organic, free range chicken ($26.99). Jack is faced with a common but peculiar quandary. His basic "psychological utility," in the sense of his ordering of outcomes in the world, are such that, if the process by which the industrial chicken was made was distasteful enough, he would prefer to spend the extra $15 on the organic chicken. However, he is not now, and likely will never be, in possession of the information required to determine whether in fact this is the case. After all, the movie is not a perfect indicator of the production process, and particularly not for the specific chickens on the shelf before him. Still, he must choose, or go hungry. The point of this article is to model Jack's incentive to believe that the conventional chicken is not really so bad as even a balanced appraisal of the movie might suggest.

As mentioned, Jack's position is peculiar but not uncommon. It is peculiar in that it is an example of *counter-preferential* choice. Even in cases where Jack's inherently preferred chicken is the organic, he may look for ways to justify choosing the conventional. This self-justification is what I call *wishful thinking*. Jack comes to believe that the moral cost of conventional chicken is less than its monetary benefit through some motivated process because he would prefer that to be true. The message of the model is that he only has an incentive to do this when he knows that he will never find out whether it really is true or not.

The intuition behind the result is that in such cases, final beliefs are based not only on observation, but on a process of *interpretation*. Any true state evokes utility results. When the state is unknown, the signals received are interpreted, a process I model as a Bayesian updating of beliefs. If the signal received is noisy, but more likely to occur in state B than state A, the shift in posterior beliefs towards B models the interpretation of the evidence that the state is "more probably B" than previously believed. Note, however, that as the decision-maker doesn't learn whether it really is A or B for sure, the "probably" of the estimate (interpretation) is weighted by the prior beliefs. Therefore, as long as the information remains imperfect, the decision-maker (like Jack) can reverse-engineer a favorable interpretation of the signals by adopting strong priors in the favored direction.

Jack's quandary is common because utility, in the psychologically deliberative sense of factors that influence decision making because they "matter" to the decision maker, often includes components that go beyond materially observable facts. For instance, many people are motivated by a desire to conform to standards of ethics or fairness; such standards are necessarily unobservable in an absolute sense, and often in any kind of quantifiable sense. While it may be clear that one action is more "responsible" or "ethical" than another, the "marginal rate of substitution" between ethics and, say, money is fundamentally different from that between, say, apricots and cherries. This is not a question of refinement of "moral sensibility." A decision-maker may well know how much she is willing to compromise her moral code for a dollar, but even that knowledge, precise as it may be, does necessarily tell her how much any particular lie has done so, as this may rely on many factors that the decision-maker will never know.

Notice this is different from the uncertain quality of a particular cherry or apricot before you eat it. Any uncertainty may lead you to actions that seem *ex post* to have been mistakes. But when outcomes are noisy signals of the true state of the world, even *ex post* you cannot tell whether your choice was a mistake or not.[1]

This issue is of increasing interest in the economics literature. In a recent paper, Bénabou and Tirole (2106) describe a large body of theoretical and experimental results consistent with it. Indeed, their model is in many ways quite similar that described below. Models of cognitive dissonance reduction (Akerlof and Dickens, 1982), psychological decision-making (e.g., Caplin and Leahy, 2001), self-signaling or dual-self reasoning (e.g., Bodner and Prelec 2003) or savoring (e.g. Bénabou and Tirole, 2010) are also close to the current perspective, as is Brunnermeier and Parker (2005). Fudenberg and Levine (2006) present a model of dual-self based self deception that works on the level of preferences rather than beliefs. Mijović-Prelec and Prelec (2010) extend Bodner and Prelec (2003) with a model in which actors have three levels of belief, deep beliefs, stated beliefs and experienced beliefs. There is also a body of empirical work, discussed later, that is quite consistent with the ideas of this paper. The contribution of this model and experiment to the literature is that the criterion identified here of persistent uncertainty is common to all existing work, but has been somewhat obscured in preceding work by the variety of forms that the models have taken. Whether the source of utility is savoring, anticipatory or self-signaling, the key element of the model that generates the incentive to self-deceive is the experience of utility at a moment when the true state of the world is still unknown. That is the element identified in the model below. This is useful because it shows the relevance of self-deceptive wishful thinking in the context of nearly any social preferences, where the motivating aspect of choice is a "metaphysical" argument such as the fairness or morality of the choice.

2. Theory

The basic decision problem in this paper is relatively straightforward. An agent X must make a choice, called an action a from a finite set of actions $A = [a_1, a_2, \ldots a_N]$. In addition, there is a finite set of possible states of the world[2], $\Omega = [\omega_1, \omega_2, \ldots \omega_M]$. Payoffs π map the product of actions and states to

---

[1] This interpretive element of the evaluation of the outcome of a decision formally implies that the decision is characterized by belief-dependent utility. However, the reverse is not the case. In so-called psychological games (Battigalli and Dufwenberg, 2007, 2009, building from Geanokoplos, Pearce and Stracchetti 1989) preferences depend on beliefs, and outcomes are interpreted in a similar manner. However, there is no indeterminacy. The interpretation is clarified by the equilibrium structure of the outcome. By assumption in equilibrium, the interpretation of the observable outcome is determined. Such, indeed, is the definition of the equilibrium. Thus in these equilibrium models, preferences do depend on beliefs, but the epistemological problem is solved by the equilibrium itself.

[2] I assume that the action space is constant across states because I will later allow X to assign positive probability to all states. If some states allowed different actions, then she could make inferences about the state by observing the possible actions. In

the real line. I further suppose the structure of the choice is a Pareto-ranked coordination "game" against Nature:

A1: For any state of the world, there is a single action that gives the highest payoff.

A2: Any action is payoff-maximizing for at most one state.

A3: States can be ordered such that the payoff earned by the maximizing action is a decreasing function of the order of the states.

The decision that *X* has to make has three steps. First ($t = 0$) she develops prior beliefs about the state of the world. These beliefs, as mentioned above, take the form of a probability distribution over states. The specific process by which I assume the priors are generated will be discussed below, but for now I underline that I suppose there is some kind of psychological commitment to these priors, so that they maintain relevance in the following steps. In the second phase ($t = 1$), *X* makes a decision that maximizes expected payoffs given these beliefs. In the third and final step ($t = 2$), she gets more information about the true state, and updates her beliefs in a Bayesian manner.

There is nothing very unusual in this set-up, unless it is the fact of breaking the decision down into so many parts. At this point, I introduce two new assumptions. The first is the problem that I am interested in investigating: self-deception. In phase $t = 0$, the agent begins with exogenous information that leads to non-degenerate prior beliefs $p = \{p_1, p_2, \ldots, p_M\}$. However, she may costlessly adjust these beliefs to some $\rho = \{\rho_1, \rho_2, \ldots, \rho_M\}$, which will then become her "committed" beliefs. Thus, I make the assumption that

A4 (self-deception): For any state $\omega \in \Omega$, ρ_ω may not be equal to p_ω, although[3] $\sum_{\omega=1}^{M} \rho_\omega = 1$

I abstract from any mechanism by which this distortion process may take place. My question in this paper is not how self-deception occurs, but why, or in what cases. However, it may strike many readers as jarring that the agent appears therefore to be able to "ignore" some, but not all information. In effect, the *p* with which she is endowed when she approaches the problem can be manipulated in the model at time $t = 0$, but the additional information she observes at $t = 2$ is fully incorporated through the constraint of Bayesian updating. While I leave a rigorous explanation for this difference to future work, I would suggest a possible informal explanation consistent with the principles, for instance, in Bénabou and Tirole (2016). The objective, initial priors *p* may often be induced from a

effect, Ω refers to all the possible states given the actions available – all those that cannot be ruled out. Compare the discussion in Yildiz (2004).

large number of similar, but not identical cases to the one at issue in the current decision problem. They may therefore take the form of “general information”, whose basic trope is well understood (lying is wrong; honoring commitments is right; equal divisions are fair; one good turn deserves another, etc.), but which lacks clear detail related to the current situation. The information in p is in some way “vague”. On the other hand, the information, if any, which arises following the decision is directly and clearly linked to the current decision context. It is very specific to the choice that the agent has just made, and is therefore “vivid”. Thus my informal explanation for why it may not be a wholly unreasonable modeling choice to allow the agent to manipulate p, but not the information that comes as a result of the information turns on the idea that vivid information is harder to ignore or mentally distort than vague information.

The second assumption provides the solution I offer as a necessary and sufficient condition to make such self-deception rational. In phase $t = 2$, the additional information that X receives may not be fully revealing. Specifically, X observes a signal q from some set Θ that maps to probability distributions over the states. This mapping shapes X's interpretation of the signal as suggesting a particular state. Specifically, it is known which signal $q(\omega)$ is most likely to occur in state ω for any ω in Ω. For simplicity, suppose that there is a “better than average chance” of getting the “right” signal, and an equal, “worse than average” chance of getting any signal other than the right one.

A5 (imperfect revelation): For any x in Ω,

$$\Pr\left[q = q(x) \middle| \omega = x\right] = \theta, 1 \geq \theta > \frac{1}{M}; \Pr\left[q = q(x) \middle| \omega \neq x\right] = \varepsilon = \frac{1-\theta}{M-1}.$$

Having observed this signal, X updates her beliefs by Bayesian reasoning, arriving at final beliefs μ, defined as follows:

$$\forall x, y \in \Omega, \mu_x = \Pr\left[\omega = x \middle| q(y)\right] = \begin{cases} \dfrac{\rho_x \theta}{\rho' q(y)} & \text{if } x = y \\ \dfrac{\rho_x \varepsilon}{\rho' q(y)} & \text{otherwise} \end{cases} \tag{1}$$

where the denominator is a vector multiplication. Given these assumptions, X's decision problem is to choose ρ in $t = 0$ to maximize expected utility in $t = 2$. Crucially, in $t = 0$, X knows that the choice of ρ will determine her action in $t = 1$, but cannot have any effect on the true state of the world or the likelihood of seeing different signals. The question becomes: When will X choose $\rho \neq p$? The answer is the main proposition of the theory:

[3] The restriction to additive distorted beliefs can possibly be relaxed. Exploring the relation that doing so would create between self-deception and ambiguity is an interesting avenue for future research.

Proposition: $\rho = p$ is optimal if and only if $\theta = 1$.

Before demonstrating this proposition, note that from the perspective of $t = 0$, the expected utility yielded in $t = 2$ by action a given potential committed beliefs ρ is

$$EV(a|\rho) = \sum_{s=1}^{M} p_s \left(\theta^2 \rho_s \pi(a,s) + (1-\theta)\varepsilon \sum_{x \neq s} \rho_x \pi(a,x)\right) \frac{1}{\rho' q(y)} \tag{2}$$

In this expression, the exterior sum is over p, while the interior sums incorporate ρ. This difference reflects the fact that the real world is not affected by the self-deception, but the actions taken are, and crucially, the interpretation of the outcome may be. Next, the coordination-game structure of the decision implies that for any state ω, there will be a threshold probability ρ_ω^* such that, given a distribution, if the corresponding component of ρ is greater than ρ_ω^*, X will choose the action that maximizes payoff in ω.[4] Call this action $a^*(\omega)$. Now we proceed with the proof.

IF: Suppose $\theta = 1$. This case is straightforward expected utility maximization. Both the bottom branch of (1) and the second term inside (2) disappear.[5] Further, the top branch of (1) is identically equal to 1. This implies that expected utility does not depend directly on ρ in $t = 2$, which again reflects the irrelevance of the distorted beliefs to real events. The expected utility for any action is simply $\sum_{s=1}^{M} p_s \pi(a,s)$, which – given the assumed structure of the problem – implies an almost single-valued[6] optimal response correspondence of beliefs: the $a^*(\omega)$ for whatever state is favored by p and $\pi(.)$. However, the action taken at $t = 1$ does depend on ρ. In finite state spaces there will generally be many values of ρ which lead to the optimal solution. For instance, any $\rho_\omega \geq \rho_\omega^*$ will do so. Since by construction p_ω leads to $a^*(\omega)$, it must be that $p_\omega \geq \rho_\omega^*$. Therefore $\rho = p$ is optimal. ONLY IF: Now suppose $\theta < 1$. Focus on the actions that maximize utility in some state. Expression (2) implies that the expected utility of each of those actions will be monotonically increasing in the prior likelihood laid on the state in which it maximizes utility. This in turn implies that the maximum expected utility must be a corner solution in ρ; the optimum beliefs are degenerate. Since the exogenous priors p were non-degenerate, this implies that at the optimal solution, $\rho \neq p$. *QED*

---

[4] This probability will be a function of the distribution across the other states, but that is not key to my argument. Because I am not putting restrictions on the kind of manipulation that can be effected on beliefs, this condition is not as strong an assumption as it might appear. It shows that there is a way to engineer any particular choice as "rational".

[5] Interestingly, this would not necessarily be true in a case of ambiguity, where ρ consisted of non-additive "capacities" rather than probabilities.

[6] It is a function, if we abstract from the countable points of indifference.

2.1. An example

Agent X lives for 3 periods t = 0, 1, 2. There are two states of the world, L or H. In period 0 her non-degenerate, objective (but vague) priors, denoted p_L and $(1 - p_L)$, are determined by some exogenous process. At this point, by "some mechanism" she costlessly adjusts these priors to ρ_L and $(1 - \rho_L)$. She may, but need not, choose $\rho_\mathrm{L} = p_\mathrm{L}$. If she does not, she self-deceives. In period 1 she must take an action U or D to maximize her expected period-2 utility, based on the values in Table 1, below, and ρ.

State Action	L	H
U	1	0
D	0	2

Table 1: Payoffs by state

In period 2, the terminal stage, X observes a signal $\sigma \in \{l, h\}$, which is correlated with the state of the world. Specifically, $\Pr[h \mid H] = \Pr[l \mid L] = \theta$, with $1 \geq \theta > 0.5$. The parameter θ represents the informativeness of the signal; I distinguish cases where $\theta = 1$ from those where $\theta < 1$. In the latter case, I allow utility to be based on a (non-degenerate) probabilistic assessment of the situation. Define X's beliefs at the end of period 2 as μ_σ, $(1 - \mu_\sigma)$ that the state of the world is L and H, respectively, based on the signal $\sigma = h, l$. The decision tree for this scenario is shown in Figure 1, below.

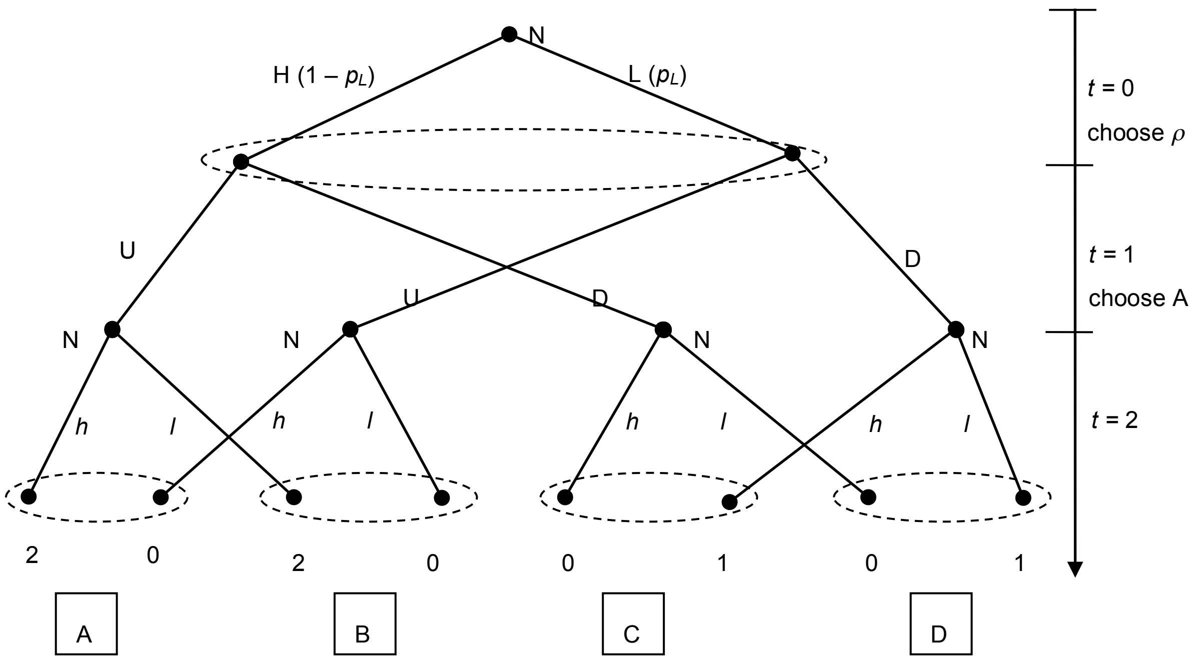

Figure 1: extensive form of the decision context.

Again, at $t = 0$, X chooses beliefs ρ, knowing that while these will determine her choice at $t = 1$, they can have no impact on the true state of the world, or the probability of seeing signal h or l. Thus in period 0, X will use the objective prior, p, and the informativeness of the signal, θ, to determine what beliefs, ρ, will yield the greatest expected utility in period 2.

Given the values in Table 1, X knows in period 0 that in period 1 she will choose action D if and only if her beliefs are $\rho < 2/3$. The final beliefs, which determine utility, are simply the Bayesian posterior probabilities given the signal θ and the chosen priors ρ. Upon seeing a low signal l, the belief that the state is L is

$$\mu_l = \Pr\left[L \middle| l\right] = \frac{\rho\theta}{\rho\theta + (1-\rho)(1-\theta)} \tag{3}$$

On the other hand, if the signal h appears, X will believe she is in state L with probability

$$\mu_h = \Pr\left(L \middle| h\right) = \frac{\rho(1-\theta)}{\rho(1-\theta) + (1-\rho)\theta} \tag{4}$$

In this example it is clear that if $\theta = 1$ then (3) equals 1 and (4) equals zero, for any ρ. This breaks the information sets A, B, C and D in Figure 1, and implies that the expected value of μ is p, regardless of ρ. Suppose without loss of generality that $p < 2/3$, so based on p, it is rational for X to choose D. Would it be rational for X to choose to adjust her beliefs? Any belief $\rho < 2/3$ will result in the same action as p. Hence, from the pre-adjustment point of view, there would be no incentive to choose any $\rho < 2/3$ different from p. An adjustment to $\rho > 2/3$ would result in choosing U instead of D at $t = 1$. Given p, this would result in a lower expected payoff, so X does strictly better when $\rho = p$ than for any $\rho > 2/3$. Thus when $\theta = 1$, the exogenous beliefs are optimal.

On the other hand, if $\theta < 1$, then the expected utility for each available action A = U, D is

$$\begin{aligned} EV\left(A \middle| \rho\right) = \Pr(l)\left[\mu_l U(A,L) + (1-\mu_h)U(A,H)\right] \\ + \Pr(h)\left[\mu_h U(A,L) + (1-\mu_l)U(A,H)\right] \end{aligned} \tag{5}$$

For any action, and any value of θ in [0.5,1), this function is monotone in ρ. Therefore, the optimal beliefs will be corner solutions. Since by assumption the priors are non-degenerate, this means that the optimal beliefs are different than, and in particular, more extreme than the objective priors for all values of θ in [0.5, 1). Furthermore, the expected utility will be higher if X chooses to believe the state is H and choose action D than in the reverse case. Figure 2, below, shows the value functions for each action when $\theta = 0.9$ and $p = 0.5$. As θ goes to 1, the functions for actions U and D converge to horizontals at 0.5 and 1, respectively. However, for any value of $\theta < 1$, the "envelope function," defined as the expected value of the action which does best, has a global maximum at $\rho = 0$.

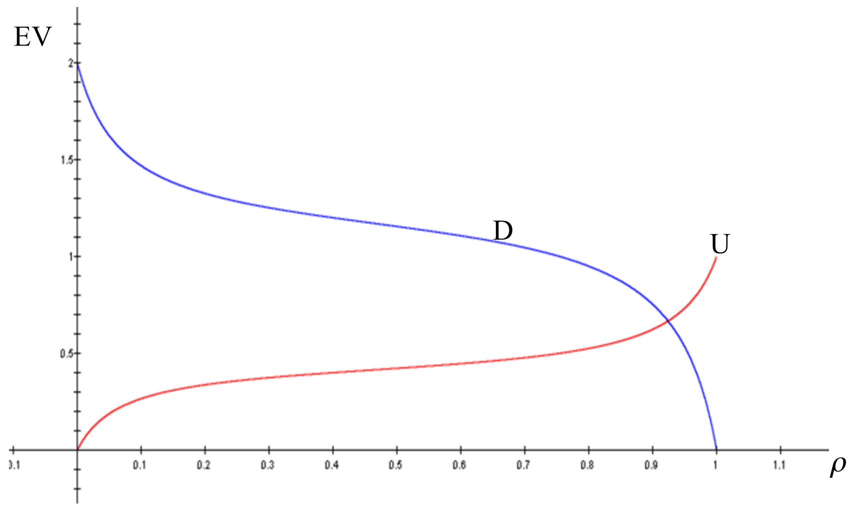

Figure 2: Value functions of actions *U* and *D* by ρ when $\theta = 0.9, p = 0.5$

2.2. Discussion

This phenomenon of, in some way, deeply preferring one option, but choosing another, cannot be directly analysed in the standard choice-theoretic framework described, for example, by Anscombe and Aumann (1963). On the one hand it treats different kinds of events. The basic premise of a lottery in that framework, whether it is a "roulette wheel" type of risk-flavored uncertainty or a "horse lottery" of more ambiguous nature, or some compound combination of them, is a device that determines which of several *a priori* uncertain events actually obtains. That determination is precisely what is called into question in cases such as Jack and the chicken above. In lotteries of the type normally considered in decision theory, the outcome can be thought of as a perfectly revealing signal of the true state of the world. That is, they describe situations where the *ex post* beliefs about the state are concentrated on one possible option, at least along all the dimensions that matter to the decision-maker. You see which case the roulette ball has landed in; you know after the race is over which horse has won. These are by assumption the only attributes of the initial decision problem that are meaningful or important to the decision-maker, thus all payoff-relevant uncertainty has been revealed. But these are the limiting case of all possible signals about the true state. In a more general framework, the information you get may not perfectly reveal the outcome. Imagine watching the horse-race on television, and that the reception is suddenly lost during the last lap of the race. You know where your horse was when the picture went black, but that is an imperfect signal of which horse actually won. The other extreme case, exemplified for instance by Jack and the Chicken, is one where the "outcome" actually reveals no information at all about the true state. As another example, consider a poetry contest[7]. Say I submit a poem, and it does not win. This is a signal that maybe I am not a great poet, but it is not a perfect signal, since there are many reasons other than my limited poetic ability that might explain my not winning any particular contest. Winning will (one hopes!) be correlated with poetic ability, but the space of possible true states after not winning any finite sequence of contests still conceivably contains the state of "unappreciated genius".

The nature of the uncertain phenomena in standard decision theory is too restrictive for an appropriate treatment of wishful thinking. But there is another, in some ways more fundamental incompatibility that comes directly from the counter-preferential nature of the choice when behavior is shaped by wishful thinking. In the standard theory of decision under uncertainty, preferences are over acts. According to the most stringent definitions, these preferences have no "real" ontological validity, but are purely analytical tools of the observer, defined by the choices. If we accept that revealed preferences are essentially an analytic tool to describe choice, whose relation to the psychological experience of the chooser is more or less irrelevant, then counter-preferential choice is in fact contradiction in terms, along the lines of a square circle, or a parent who has never had children. Thus the formal model of preferences over acts in standard decision theory rules out wishful thinking by definition. More precisely, it defines its object outside the scope of wishful thinking. The "true" beliefs are by definition those that shape behavior. Note that in this case, the axioms of choice under uncertainty simply describe the conditions under which the expected utility framework can be used to describe choice, so the conclusion goes no farther than the position that expected utility cannot be used when people are wishful thinkers in this sense. We will see later the limitation of this position. The model presented essentially extends the domain of acts onto the beliefs themselves. This results in not all acts being observable, so the link between observable choices and revealed preferences is broken, but in the larger domain it does not violate the axioms of the theory, and so expected utility can still be applied to that wider domain, even though it is harder to pin down preferences from (the observed part of) chosen acts.

This can be seen again if we consider the more intuitive view of the rational choice under uncertainty paradigm, in which choices are considered to reveal some real, but otherwise unobservable beliefs or preferences. If we follow this view to admit the psychological reality of preferences and beliefs, as most treatments do at least in the motivation of the formal theory, we claim that choice theory axioms delineate the cases when observed choices can be taken to reveal these preferences. Here we find somewhat the reverse problem with wishful thinking to the above. If wishful thinking is behaviorally relevant at all, then its very definition is the process by which people who prefer option x nevertheless

[7] Thanks to Claude Fluet for this example.

justify choosing y. Thus a study of wishful thinking appears to exclude precisely those cases that the standard theory of choice under uncertainty describes.

But while decision theory in its complete form is not amenable to description of self-deception, returning to a utility framework allows a more precise description of the phenomenon. While decision theory is developed in a formal structure that does not explicitly rely on utility, the structure of the acts that it treats is discussed as being composed of outcomes that receive utility numbers, and beliefs that map outcomes to states. Thus the fundamental vision of human experience one gets from standard decision theory is utilitarian. Maintaining this vision, it follows that the process by which people change their act-preferences should concern at least one of these components: utility or preferences. Wishful thinking is either a manipulation of the utilities that outcomes seem to be "worth", or the beliefs that link the acts to the expected outcomes. In this paper, I consider the latter possibility.

Beliefs, which refer to the standard concept of a subjective probability distribution over the conceivable states of the world, are ordinarily considered to be the information the decision maker uses as a guide from actions to outcomes, and are implicitly or explicitly assumed to be exogenous. So far as they fill this role, it is relatively immediate that objectively true beliefs – for instance, ones which correspond to the objective relative frequency of different states in some long run – are "preferable". If I get \$5 for correctly calling a coin toss, and lose \$1 for missing it, then in general I would prefer my beliefs to correspond to the true probabilities involved when deciding on how much to pay for the game. But beliefs can also have an affective role. Consider flipping the coin n times, and my beliefs about flip number $n + 1$. In this case, since I will never know the true value, the incentives change, although the direction of the change is not obvious. One could imagine a "sour grapes effect", in which I would prefer to believe I would have missed that toss. Alternatively, one could imagine that I might prefer to believe it would have saved me.

3. Experiment

3.1. Design

Empirical testing of this theory faces a significant hurdle. Simply eliciting beliefs and behaviour is insufficient, as the model shares with standard rationality the assumption that beliefs should be aligned with behaviour. The line of causality in the model runs counter to that in standard theory – you do not act in function of exogenous beliefs, but rather choose what to believe as a function of what you want to do – but the action-belief pair should be consistent on both counts. Also, experimental manipulations that affect incentives to believe in one state or another will usually also affect incentives for choice, so independent manipulation of the decision context and the belief incentives is challenging. However, a variable that is identified in the model, which determines the extent of belief distortion, is the informativeness of the signal. The question is how to vary this informativeness without at the same time influencing the incentives of the decision itself. The strategy adopted here is that of choosing whether or not to play a particular bet.

Participants in this study were presented with binary lotteries. Figure 3, below illustrates the form. Each lottery was composed of a quantity of money, and a probability of winning it. They were informed that at the same time another lottery had simultaneously been generated. They had the options either to play the current lottery, or to pass it up and take the second; but if they passed on the current lottery, they would not be able to go back to play it after seeing the other. Thus, there is a single, binary relevant choice: pass or play the current lottery.

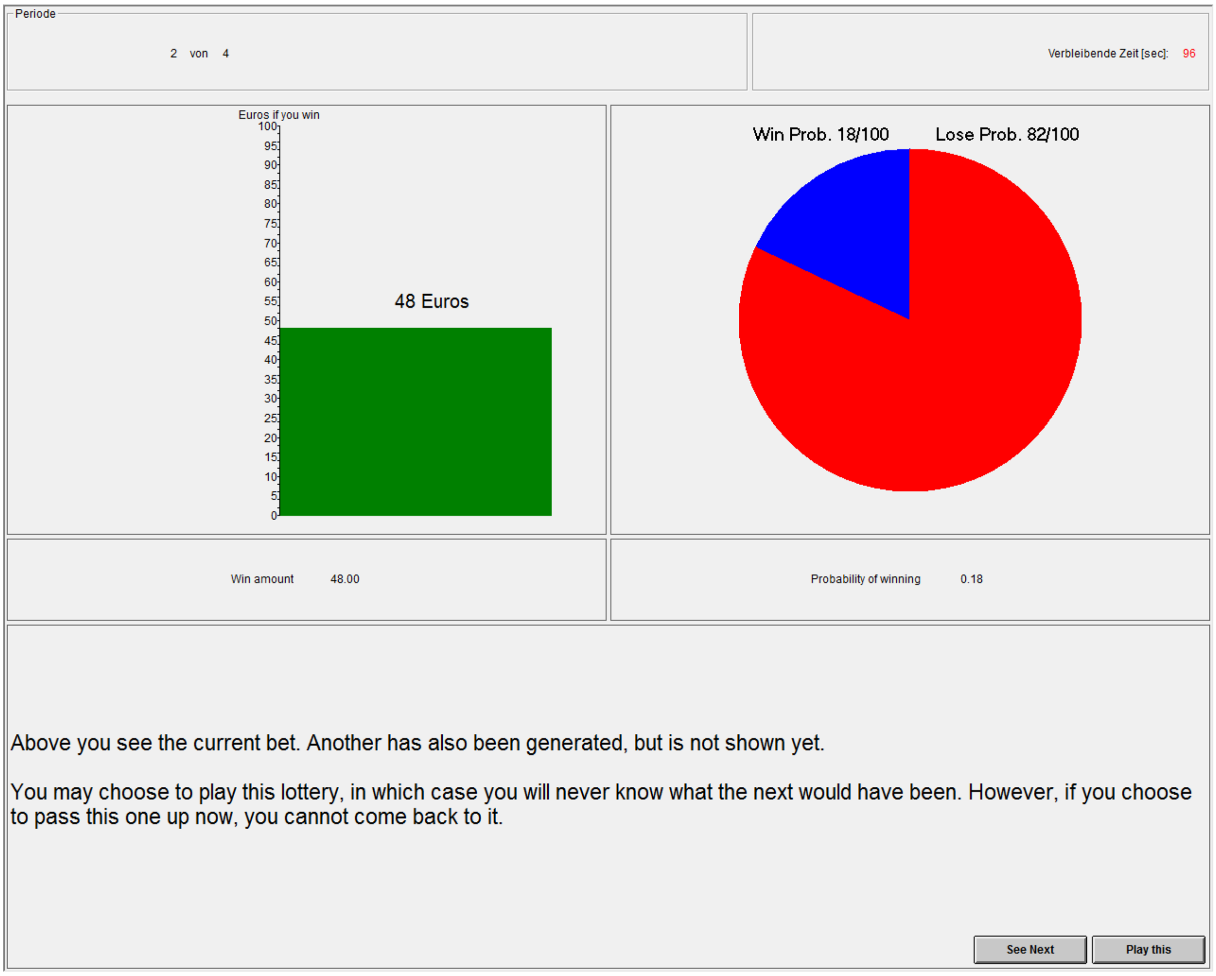

Figure 3: screenshot of the experimental interface

Participants face a sequence of 20 randomly generated choices of this kind. After making the choice, they are shown the realization of a random event, framed as a bingo ball with a number between 1 and 100. If the number is lower than the probability of winning the bet, they earn the amount indicated. They were not given any information about the distribution of the lotteries, other than the sequence they observed. In fact, the expected value of each lottery was independently generated to be equal to €$(4 + 3\varepsilon)$, where ε distributed as the square of a uniform random variable between 0 and 1. Therefore the highest-value bets had an expected value near €7, while the lowest were near €4, but the mass of the distribution was greater at the low end of the spectrum. Then a probability of winning was independently generated between 0.10 and 1.0, and the win amount was calculated in consequence. Therefore prizes and probabilities were highly correlated, but probabilities were independent of expected values.

The between-subjects experimental control consisted in whether or not participants were informed of what the second bet would have been, after deciding to play the first. In the INFO treatment, they were told that they would see not only the hidden bet, but also whether or not that bet would have been won, even if they chose to play the current one. In the NO-INFO treatment, they were told they would see the hidden lottery only if they chose to play it.

In the INFO treatment, two bets were shown at the end of the period, the one chosen and the one passed up. Both bets were subjected to the same bingo ball resolution, and so participants in that treatment were informed both of what they won, and what they would have won, had they made the other choice. In the NO-INFO treatment, they were informed only of the bet they actually chose, although those who chose to *Pass* on the original lottery were presumably able to infer what their earnings would have been had they chosen to *Play* it instead.

3.2. Predictions

The first behavioural assumption is that participants will compare each bet with an expected reference bet they have in mind. They will Play any bet that they prefer to this expected alternative, and Pass on any bet that is worse.

In the terminology of the model, the information they have about the alternatives is p, the objective belief, which they may distort to some ρ before choosing to Pass or Play. When they know that the alternative bet will be shown later, they effectively get a revealing signal about the true state of the world, and according to the predictions of the theory, they should not distort their beliefs. On the other hand, when they know that they will never be faced with the alternative, there is an asymmetry of incentives across the two actions. If they Pass the current lottery, then they will have seen both bets, so they will know the state of the world with certainty. But if they Play the current lottery in the NO-INFO treatment, uncertainty over the state of the world, in terms of the optimality of one choice or the other, is preserved. Therefore they have an incentive to distort beliefs, according to the theory described above, the NO-INFO treatment towards Playing the current lottery. The formal hypothesis is

H1: Participants will choose to Play more in the NO-INFO treatment than in the INFO treatment.

3.3. Data

Four sessions were run in October, 2017 as incentivized activities in the context of an Organizational Behavior course at the Burgundy School of Business in Dijon, France. In all, there were 131 participants, 67 in the NO-INFO treatment and 64 in the INFO treatment. Participants were paid for one of the 20 lotteries they chose, determined randomly at the end of the experiment. The sessions lasted less than 30 minutes. The average pay was quite low (around €2.60), because many participants were randomly attributed a bet they did not win. However, all participants seemed highly engaged with the task, and the fact that the experiment was run during class time implies that there was a low opportunity cost to being there, so transportation costs needed not be covered.

3.3.1. Differences across individuals

Because each participant made 20 choices, the data contain 2620 individual observations, in a panel of 131 independent individuals. Because we have 20 observations for each one, we can control for individual risk attitudes in the data. To do this, I ran a separate regression for each individual participant, regressing the probability and expected value of the bet presented on the action choice in a LPM framework. The results were strikingly varied. A histogram below shows the distribution of estimated coefficients on the probability of winning the bet, defined as P-Type. As can be seen, the distribution appears to be somewhat bi-modal, with one mass near 1, indicating people who are more likely to take a bet when the probability of winning it is higher, and another just below zero, indicating those who took bets only when the prize was high enough. This is a useful measure of risk aversion, and will be used in the subsequent analysis. The E-type, a corresponding measure for the coefficient on the expected value of bets, was more normally distributed, although with a lower average (0.09 for E-type; 0.27 for P-type).

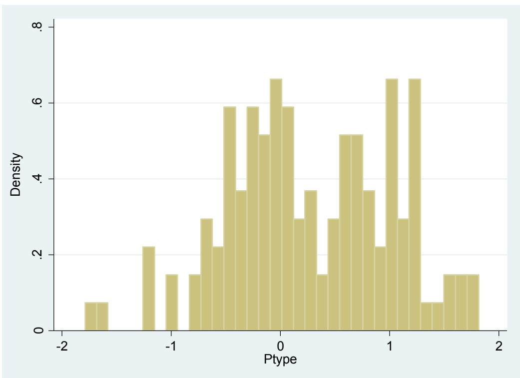

Figure 4: distribution of coefficients on probability for taking a bet

There was also a difference in the number of bets Played out of the 20 seen. The histogram in Figure 5 shows that about half of subjects Passed on about half of the bets, but that a minority took nearly all of them, and a few people passed on nearly all. Differences in this value across experimental treatments will represent the main research result from the experiment.

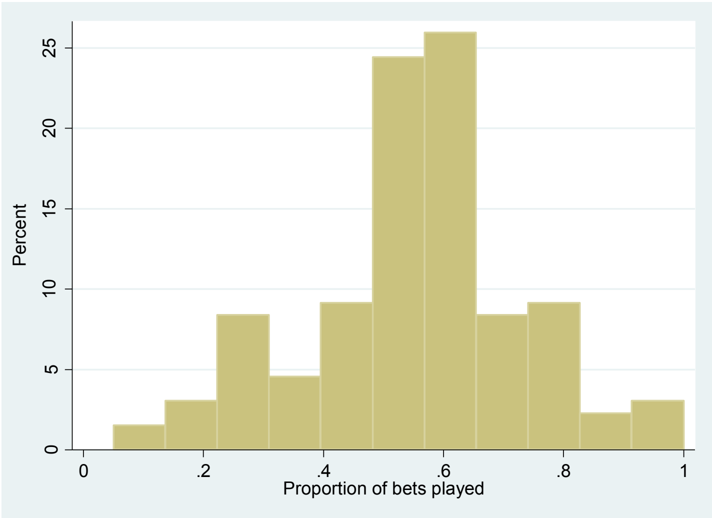

Figure 5: distribution of action choices

3.3.2. Differences across time

In addition to differences across individuals, the panel structure of the data allows us to look at changes over time. The first one the reader likely has thought of, given the design of the experiment, involves learning the distribution. One interesting metric that can be used to investigate this question is the time participants took to make their decisions. The software took a measure of the time whenever participants chose to Play or to Pass on a lottery. Therefore the measure includes the time taken to look at the outcome as well as that taken to make the decision itself. While it may have been interesting to have the more fine-grained information, for the purposes of learning about the distribution, looking at and reflecting on the outcome of a choice may have been as instructive as looking at the bet. The graph below shows the average time participants took in a period, for each period and each experimental treatment.

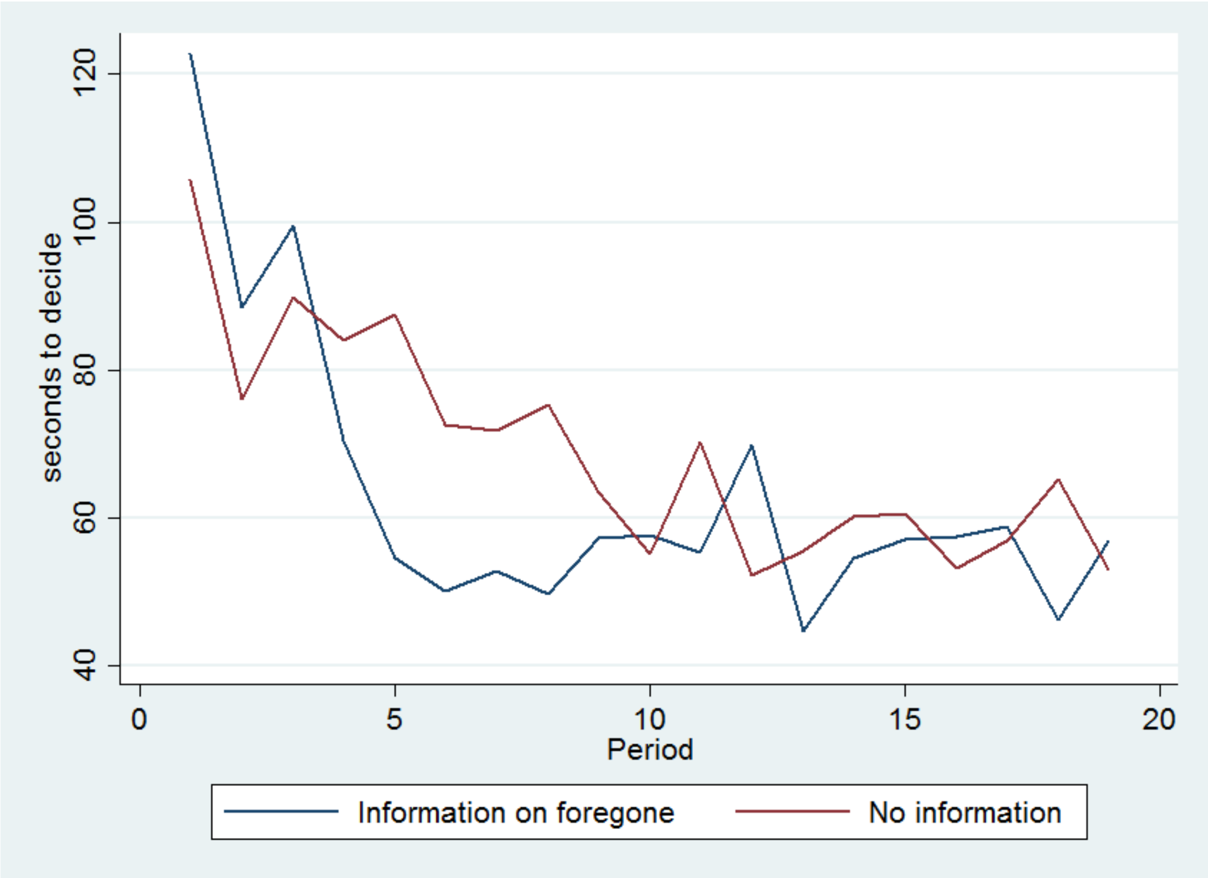

Figure 6: time between decisions over time, in each period.

Informally, there are three interesting things about this relationship. First, the time in both treatments converge to just under one minute per choice from period 10 on. Second, the time in the INFO treatment, where participants learned about both the current and foregone lotteries when they decided to play the one offered, converges to the same level, but much earlier, apparently from around period 5. This suggests the intuitive idea that those participants are getting more information about the distribution, and are therefore learning faster what constitutes a "good" lottery and what a "bad" one. Third, and related, the INFO treatment actually spends more time in the first periods than does the NO-INFO treatment. This is again suggestive that participants in the INFO treatment are attempting to integrate a larger amount of information. These patterns are also statistically significant. A regression on the time spent, limited to the first 10 periods, that interacts the experimental treatment with the period number, shows significantly positive effects for the INFO (β = 13.8, p = 0.002), implying a nearly 14-second difference in time spent in the first period, and while both groups have negative slope, a significantly negative interaction term (β = -4.6, p = 0.000), indicating that the time fell by more than four seconds more each period for those in the INFO group than in the NO-INFO.

It is not entirely clear what the appropriate way to treat this effect is in terms of the model of wishful thinking I want to test. On the one hand, the initial response is that one should wait for behaviour to converge before doing analysis, since that represents something closer to participants' "real" revealed preferences. On the other hand, the model in Section 2 above is explicitly predicated on the actors' not having perfect ideas about the state of the world. While for simplicity in the model there was no cost to belief manipulation, it seems like a plausible conjecture that the more diffuse the beliefs – which is to say the more possible things that could happen from the subjective point of view – the easier it would be to believe one or another in a motivated way. Therefore it may be an only slightly augmented model that predicts that the wishful thinking I want to study should diminish as the time converges, and in fact the part of the sequence where the phenomenon is most likely to occur is between periods 5-10, when the ideas of the INFO group are clearer than in the NO-INFO group. Still again, it could be argued that before period 5, when nobody seems to have clear ideas, the model may apply equally to both groups. None of these options, moreover should be considered as overly attenuating the model's external validity, as most people are faced with relatively new decisions quite frequently. Comparing the aggregate effects of mechanisms that operate mainly in new versus familiar choice situations is a problem for another paper, but neither seems to be premie facie uninteresting.

The idea that participants are learning about the bets as they go is again corroborated with an investigation of the kind of bets they accept over time. The graph below shows the average probability of lotteries Played and Passed in each period

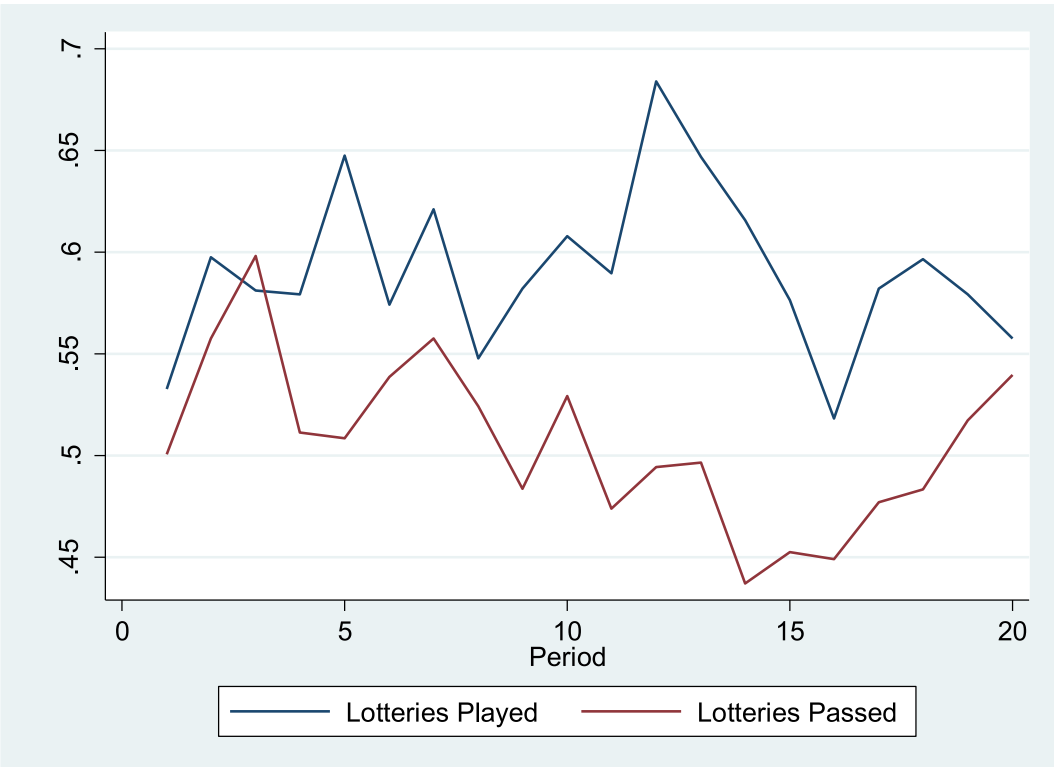

Figure 7: Average Probability Passed and Played over time

This graph shows lotteries played over time as more or less constant in probability, while those passed up seem to fall in probability. The relevant regression, regressing the probability on the interaction of period and choosing to play it, confirms this. The difference of the intercepts is marginally significant ($\beta = 0.05$, $p = 0.053$), but the number of observations has fallen to 40. Those not played to diminish significantly in probability over time ($\beta = -0.003$, $p = 0.028$), but the test of the sum of Period and interaction coefficients is not significant ($F_{1,36} = 0.01$, $p = 0.909$), indicating that the probabilities for bets taken do not change. So participants were learning to reject bets with low probabilities.

3.4. Results

The main question of interest in this experiment was whether information had any impact on participants' probability to play a given bet. The regression table below shows the effects of the treatment, clustered at the participant level, and using as controls the coefficients from the individual regressions run, which show the importance to each individual of probability and expected value on choice. There are four regressions reported, which break the effect down by the period groupings suggested above. The first column shows the overall effect, including all periods. The second shows the effect for the last half of the treatment, when speed of decisions, and hence potentially beliefs about the distribution of lotteries, has converged. The third shows the interim period during which those with information have converged to their long-run time/beliefs, while those with less information are still learning. And the last shows the effect during the first four periods, when both groups are learning.

VARIABLES	*All periods*	*Periods 11-20*	*Periods 5-10*	*Periods 1-4*
INFO	-0.0558*	-0.0428	-0.109**	-0.0406
	(0.0310)	(0.0385)	(0.0426)	(0.0430)
P-type	0.0469**	0.0407*	0.0482*	0.0507*
	(0.0185)	(0.0234)	(0.0285)	(0.0264)
E-type	-0.0948	-0.0626	-0.167	-0.120
	(0.116)	(0.140)	(0.161)	(0.173)
Constant	0.578***	0.568***	0.627***	0.572***
	(0.0295)	(0.0336)	(0.0381)	(0.0394)

Observations	2,560	1,280	640	512
R-squared	0.009	0.006	0.019	0.009
Robust standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1				

Table 2: Regression results of the treatment effect on behaviour.

The results show that in all specifications, information decreases the probability that a bet is played. There is a marginal effect overall (p = 0.07), which is significant only in the segment of the experiment between periods 5 and 10 (p = 0.02), which is also when the time taken in the two treatments differs most. It therefore appears as though uncertainty about the actual distribution helps the wishful thinking effect. While this is not strictly speaking a prediction of the theory, it is quite in line with the spirit of the model, as that uncertainty is what permits the “wiggle room” that wishful thinking exploits.

4. General Discussion

4.1. Relation to regret theory

After the experiments were run, a relationship – in hindsight, quite clear – to regret theory, as formalized by Loomes and Sugden (1982), emerged in the design. According to that theory, which was developed to explain decision paradoxes including the Allais paradox, utility for any outcome is adjusted by a function representing the regret or rejoicing experienced due to the negative or positive (respectively) difference between that choice and what would have been obtained from another choice. The utility function for choosing action a over action b, when each gives a payoff $x(i)$ for $i = a, b$ is therefore of the form

$$U(a) = x(a) + R\big(x(a) - x(b)\big)$$

for some convex, non-decreasing, three-times differentiable function $R(.)$ such that R(0) = 0. In the current context, action a would refer, for example, to the choice to *Play* a given lottery, called *This* below, while b would refer to the *Pass* option, opting for the *Next* lottery. To construct a benchmark case for investigating the extent to which regret theory can also explain these results, consider the following auxiliary assumptions

R1: Each bet is compared to a hypothetical *reference bet*. If *This* lottery has the form pW for winning prize W with some probability p, denote the reference bet *Next* as qM, winning prize M with probability q.

Assumption R1 abstracts for reasons of tractability from risk aversion over the distribution of possible bets. Given the relatively low values in the experiment, this does not seem unreasonable.

R2: Suppose that the expected value of the lotteries is the same, so $pW = qM$.

This assumption is a *ceteris paribus* condition that makes the analysis easier, and with continuous utility will also have the same implications for lotteries close enough in practice.

R3: Suppose that no regret is experienced when the foregone bet remains unknown

R3 is an alternative to making assumptions about how uncertain outcomes enter the regret function. Although strong, there are at least four reasons why I make it. (a) The original theory deals with known outcomes, and so introducing unrevealed outcomes departs from the theory. (b) The assumption is consistent with literature on regret, both theoretical (e.g. Zeelenberg and Pieters 2007, Humphrey 2004) and empirical (e.g. Hetts et al., 2000; Ritov, 1996; Larrick and Boles 1995). It is also consistent with principles of information avoidance (e.g., Golman, Hagmann and Loewenstein, 2017) – what you don't know will hurt you less (c) It is a tractable way to deal with the model. (d) A *priori* it seems likely to increase the chances that regret aversion can also explain the experimental results, and so it is a conservative strategy for my purposes.

Given this, the observed pattern of behaviour (a higher likelihood to take a bet when the foregone outcome remains unknown upon doing so than when it is revealed) turns on the sign of the anticipated regret choosing to *Play This*. Figure 8 below illustrates. V(*X*) refers to the value of the lottery, which for the current purposes I take to be its expected value.

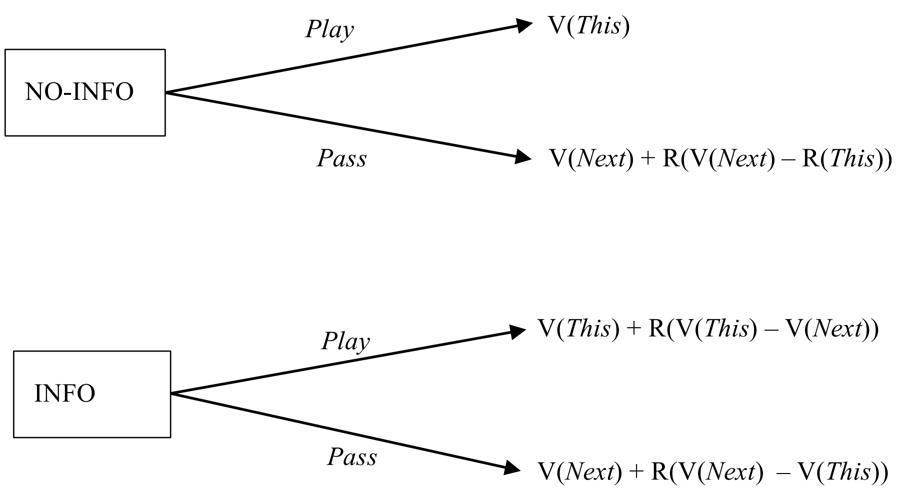

Figure 8: Experimental treatment under regret.

The NO-INFO treatment eliminates (by R3) the regret from choosing to *Play*, but not from choosing to *Pass*. Therefore, regret theory also predicts higher rates of *Play* in the NO-INFO treatment if the expected regret from choosing *Play* is negative, that is, if the average value of R(V(*This*) – V(*Next*)) < 0.

The structure of the resolution of the lotteries has an implication this expected value. Recall that in the INFO treatment, both bets displayed are resolved with the same random "bingo ball". Therefore, the structure of the dependence leaves three possible outcomes: both bets would have won; the bet with the higher win-probability would have won, but the other would not; and neither bet would have won. The expectation of the regret depends on the relative values of p and q. If $p > q$, then *This* lottery has a higher win probability – and therefore by R2 a lower payoff – than the *Next*. With probability q the regret of choosing to *Play* will therefore be $R(W - M) < 0$, because with that probability both lotteries would have been won. In words, this case is the regret felt when you take the safe bet and then see the risky one would have paid more. With probability $(p - q)$, the rejoicing is $R(W) > 0$. This corresponds to the rejoicing of choosing the safe bet when the risky one would have lost. With probability $(1 - p)$ both bets lose, and so there is no regret. Therefore, the expected regret for choosing to *Play* bets with a probability higher than the reference bet is

$$qR(W - M) + (p - q)R(W). \qquad (6)$$

Notice first that this expression is equal to zero when $p = q$, since by R2 when $p = q$, $W = M$. This is intuitive, since in that case both bets and outcomes are identical, and so there can be no regret. For small differences – *This* bet slightly safer than the reference – the result is somewhat ambiguous, as this both increases the regret in the first term and the probability of rejoicing in the second. However, the effect is linear in the latter term, and works through the regret function in the former. Without more information on the form of $R(.)$, it is therefore difficult to reach any conclusion. As mentioned, this function is usually not more restricted than being non-decreasing, three-times differentiable and convex, with $R(0) = 0$ (Bleichrodt and Wakker, 2015). Using the fact that $W = M\,p/q$ and taking the derivative of (6) with respect to p yields

$$\frac{\partial}{\partial p}(.) = -qR'(W - M)\frac{Mq}{p^2} + R(W) - (p - q)R'(W)\frac{Mq}{p^2}$$

If (6) is sufficiently convex, then the first and third terms of this go towards zero as p approaches q, leaving only (or mainly) $R(W) > 0$. This suggests that for bets somewhat safer than the reference bet, the expected rejoicing from winning something rather than nothing outweighs the expected regret from not winning the higher reference prize.

Prediction R1 (regret theory – safer than reference bet) INFO should increase the rate of Play for bets just safer than the reference bet.

When $p < q$, that is, when the current bet looks slightly riskier than the (equal expected value) reference bet, the parallel argument implies that with probability p the rejoicing from choosing *Play* is $R(W - M) > 0$, and with probability $(q - p)$ the regret is $R(-M) < 0$. The expected regret for choosing to *Play* lotteries perceived as riskier than the reference level is therefore

$$pR(W - M) + (q - p)R(-M). \qquad (7)$$

This is also equal to zero when $p = q$. The second term increases linearly in p, while the first is multiplied by the regret function. The derivative yields

$$\frac{\partial}{\partial p}(.) = R(W - M) - pR'(W - M)\frac{Mq}{p^2} - R(-M) + (q - p)R'(-M)\frac{Mq}{p^2}$$

Again, the term $-R(-M) > 0$ will dominate as p goes towards q if R is sufficiently convex. Therefore the derivative is positive and the value is zero at $p = q$, which means that the value is negative for small differences – bets seen as slightly riskier than the reference bet. This suggests that INFO will make bets slightly riskier than the reference bet look worse than will NO-INFO.

Prediction R2 (regret theory – riskier than the reference) INFO should decrease the rate of Play for bets just riskier than the reference bet.

Putting these predictions together, we have

Prediction R (Regret Theory): The distribution of bets played in the INFO treatment should be safer than the distribution of bets played in the NO-INFO treatment.

This prediction can be quite readily tested with the data. Figures 9 shows the distribution of probabilities in the bet, conditional on the INFO or NO-INFO treatment, and on the action chosen. It is clear that there was no large effect of the treatment on the probabilities. The combined Kolmogorov-Smirnov test of equality of distributions fails to reject equality in either the bets played (p = 0.458) or passed (0.893). Thus we conclude that regret does not seem to be driving these results.

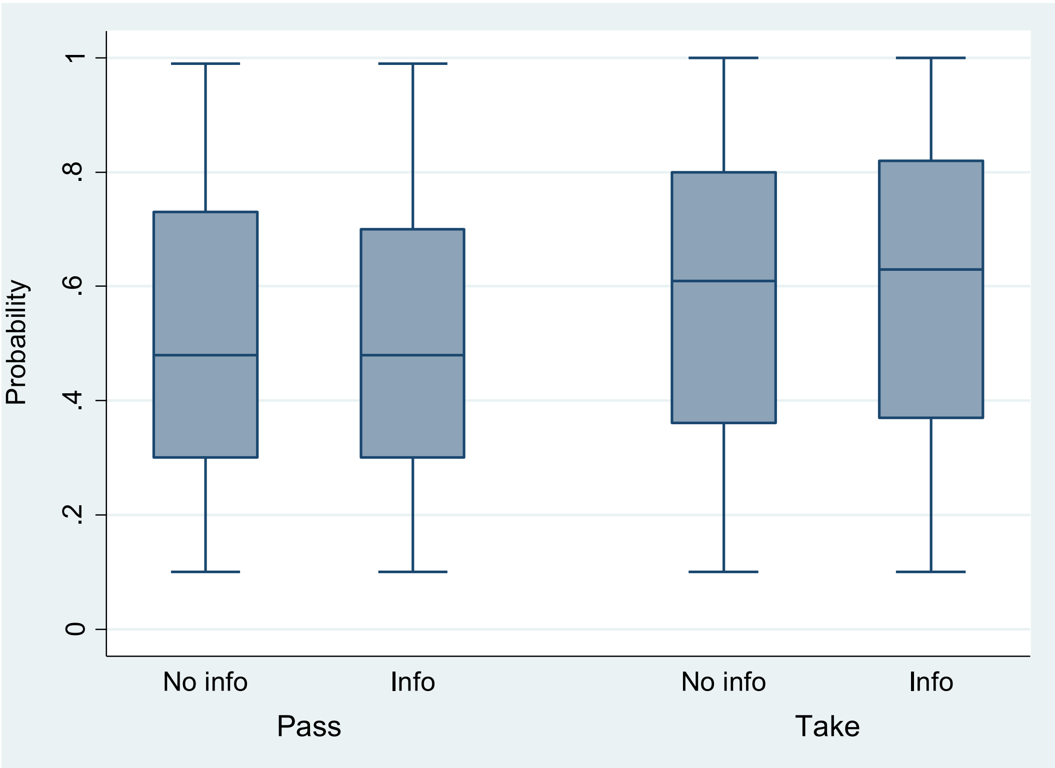

Figure 9: Box plots of distribution of Probabilities conditional on treatment and choice.

4.2. Relation to Ellsberg

Another celebrated domain that in retrospect seems related to the experimental results concerns ambiguity aversion, as illustrated in the Ellsberg Paradox (Ellsberg, 1961). In that thought experiment, agents were supposed to prefer taking a bet on an urn containing 50 percent each of red and black balls, to taking one from an urn containing an unspecified mix of balls of the two colors. Subsequent theories have rationalised this with functions that overweight the probability of extreme outcomes in ambiguous situations, making the context feel "riskier" than an equivalent one with more clearly defined probabilities. In the current experiment, as participants learn about the distribution of prizes from their experience, arguably moving from a context of ambiguity, in which they know there is uncertainty but have no clear ideas about the distribution of lotteries, to one of risk, as the distribution becomes clearer. This reduction in ambiguity would therefore naturally predict the reduction in treatment effect over time. In contrast to the link to regret theory above, moreover, this does not seem necessarily to be an alternative explanation of the results, but rather an interesting link between the wishful thinking model and the ambiguity literature. A similar trope occurs in other empirical literature, where a robust finding is that self-serving biases are more prominent in treatments that preserve ambiguity or uncertainty about the outcomes (Felson 1981; Dunning, Meyerowitz et al. 1989; Dana, Weber et al. 2007; Haisley and Weber 2008; Valdesolo and DeSteno 2008; Sloman, Fernbach et al. 2010; Feiler, 2014).

4.3. General Discussion

In cases of self-deception a person manages by some mechanism to engage in "motivated reasoning" and thereby manipulate her beliefs in some way. Roughly speaking, the self-deceiver chooses to believe x (the "target belief"), rather than any other y (each of which is logically a not-x), and specifically rather than some y which is at least as well supported by observed data as x is[8]. Such behavior raises two important questions: how, and why?

This note has addressed the second question. The first – how do people selectively integrate the available information into their beliefs – is an important empirical problem. However, granting that they can does not explain why they would choose to do so. Beliefs matter because, in general, the preferred action will change as the state of the world changes, and beliefs summarize information about the state of the world. Acting on beliefs different from those the data support seems destined to lead to sub-optimal outcomes. Yet self-deception is a very intuitive phenomenon.

The model presented above makes plain a sufficient and necessary condition for self-deception to operate as a motivation in the context of Bayesian rationality, that is, subjective expected utility maximization. The condition is an *uncertainty postulate*, in that it requires that beliefs about the true state of the world remain non-degenerate, essentially forever. Relaxing the standard assumption that states are revealed at the end of the game gives beliefs a quite radically different role. They no longer simply guide the decision-maker towards the optimal outcome. Rather, they take on affective benefit because they color the interpretation of any data which are observed to result from decisions. This interpretive character supplies the motivation to manipulate the beliefs.

---

[8] (Szabados 1973) argues that self-deception requires y is strictly better supported than x by the data; otherwise it is wishful thinking. (Mele 1997) Suggests the 4-part sufficient conditions for S to be self-deceived (1) the belief x that S acquires is false; (2) S treats data relevant to the truth of x in a motivationally biased way; (3) the biased treatment is a nondeviant cause of S's acquiring the belief x; (4) the body of data S possesses provides greater warrant for ~x than for x (p. 95).

In particular, we have found that in the case where (a) the optimal action depends on the state of the world, and (b) some states of the world are "better" than others, it is necessary and sufficient that final utilities remain uncertain for self-deception to be a motivating force. Not only that, but in these cases agents do best when they can establish "unshakeable beliefs" that their preferred state of the world is correct. Again, the intuition behind the result is that, when uncertainty is resolved, the manipulated beliefs may still affect the action choice at $t = 1$, but can no longer have any impact on beliefs at $t = 2$. The final beliefs therefore become irrelevant for the decision made at $t = 0$. By contrast, when the epistemological indeterminacy remains, ρ has a direct impact on utility through the beliefs.

The link to belief-dependent utility shows in the interpretation of expression (5) above as a utility function. This utility is maximized by choosing ρ, which then nails down μ, and thus also A. This modification has no effect on the proposition above. It makes the model more familiar, in that the "outcomes" (including the belief) are known with certainty at the end of the decision problem. However, it obscures the mechanism by which the result is obtained. The reason why $\rho = p$ is optimal if and only if $\theta = 1$ is that in the contrary case the state of the world is both relevant to the actor's choice problem and also unknown. The interpretation of (5) as a utility function also makes a more complex theoretical structure, at least graphically. Sacrificing both "praxeological realism" and theoretical simplicity to make the model conform to decision problems that use beliefs in a fundamentally different way seems unwarranted.

While the condition of persistent uncertainty is unfamiliar in the literature, moreover, in the background it does characterize all of the theoretical papers I am aware of, and I have not been able to construct an example that convincingly contradicts it. An interesting implication of this unresolved uncertainty is that it implies that the agent must have preferences that go beyond her own material well-being. Material well-being does not exist probabilistically, and hence no material benefit can possibly be perceived as long as uncertainty about its magnitude remains. *Homo economicus*, whose preferences are identical to material well-being, is necessarily impervious to such temptation. Perhaps unfortunately for the rest of us, "extra-material" incentives yield abundant opportunities. As illustrated in the introductory paragraph, one area of particular interest may be the demand for environmental goods. This is interesting because the self-deception may, in general, push people in several different directions. Those who, for instance, buy hybrid cars will downplay the environmental costs of battery disposal. Those who buy conventional chicken will downplay the benefits of ecologically sensitive production methods. There are also several examples from the previous literature. In Bénabou and Tirole (2010), agents want both to maximize their lifetime income (material incentive) and to have a good impression of that income along the way (psychological incentive). Lord, Ross et al. (Lord, Ross et al. 1979), found that people want both to have correct beliefs (on the appropriateness of capital punishment, in this example), and also that those beliefs agree with their previous ideas or general ideological framework. In studies such as (Quattrone and Tversky 1984; Ditto, Pizarro et al. 2009) subjects would like to know whether they have a proclivity to some disease, but also would like the result to be negative. Knowing their true proclivity would potentially allow people to take preventative – or at least preparative – measures, reducing the material cost the disease imposes. However, that would cost them the belief that they were healthy in the interim. Without the interim incentive, there would be no reason to manipulate beliefs in the way those researchers interpret their subjects as doing.

5. Conclusion

According to the theoretical results in section 2 above, the condition of persistent uncertainty answers the why of self-deception. Self-deception occurs because when uncertainty is unresolved, prior beliefs no longer have solely instrumental value, pushing the actor towards the best possible outcome. They also have direct benefit, since they influence the interpretation of later information. This principle seems to unify much fascinating work in economics. Savoring, procrastination, self-image management, fairness and reciprocity are all (non-exhaustive) examples of cases where an agent may have an interest in maintaining certain beliefs about choices. They all also share the characteristic that the object of the beliefs is not known with certainty at the time utility is perceived.

The model also suggests answers for the first question of how beliefs are manipulated. Essentially, X "reverse engineers" her prior beliefs to bring observed data into line with what she would prefer to be true. Knowing that the action D will give her a higher utility in state H than the action U will in state L, she manages to believe that H is likely enough to justify that choice by conditioning the observed information with on a distorted (e.g., selectively attended) base understanding, or prior. However, the essential message of the model will hold with other kinds of manipulation. For instance, X could equivalently manipulate the signal σ, for instance by interpreting it as implying a different θ. Which model is more accurate is an empirical question, and may well depend on the case considered. It remains a promising area for future research.

The experimental results presented are similar in their contribution. The experiment itself has some novelty, for instance in that it proposes a measure of risk aversion that does not require participants to fill out a list of prices. In terms of its ability to address the theoretical issues, while the theory presented in section 2 organizes the data well, we saw in section 4 that regret theory also has significant explanatory power, although the patterns of behavior over time are harder to fit. It is intuitive that learning about the distribution, as it fixes ideas about what could happen, would reduce the ability of participants to shape their beliefs. Indeed, this seems to be a finding that is common in research about the effect of ambiguity on self-serving beliefs. On the other hand, if anything stronger beliefs about foregone alternatives should increase the regret that participants feel. If this is not the case, moreover, then the two models become somewhat closer together. Indeed, the difference between regret aversion and wishful thinking may not be that great, modeling choices aside. There remains the difference that regret theory is conceived as "experiential", something that acts outside of rational, deliberative choice, while wishful thinking as modeled is part of the maximization process. However, that maximization could itself be semi-conscious, representing an automatic evaluation distortion. The paradox of wishful thinking is the ability to both believe one thing and recognize evidence that it is not true; an "unconscious wishful thinking" that acts on the same level of rationality as regret could help to undo that paradox. Furthermore, consider the basic effect on the opportunity cost of choice in both models. According to regret theory, the chosen alternative "feels worse" relative to the foregone depending on the outcome; in wishful thinking, the foregone alternative is "depreciated" relative to that chosen. They may be a constant's distance apart, and identical in terms of the effect on the perceived opportunity cost of choice. Further empirical work to distinguish which element in the pair is subject to this manipulation could therefore be useful.

Works cited

Akerlof, G. & W. T. Dickens, (1982), The Economic Consequences of Cognitive Dissonance, *American Economic Review*, 72(3), p. 307-19,

Bénabou, R., & J. Tirole, (2011). Identity, morals, and taboos: Beliefs as assets. *The Quarterly Journal of Economics*, 126(2), 805-855.

Bénabou, R., & J. Tirole. (2016). Mindful Economics: The Production, Consumption, and Value of Beliefs. *Journal of Economic Perspectives*, 30(3): 141-64.

Bodner, R., & D. Prelec, (2003). Self-signaling and diagnostic utility in everyday decision making. *The psychology of economic decisions,* 1, 105-26.

Brunnermeier, M. & J. Parker (2005) Optimal Expectations. *American Economic Review*. 95(4): pp 1092-1118

Caplin, A., & J. Leahy, (2001). Psychological Expected Utility Theory and Anticipatory Feelings. *The Quarterly Journal of Economics*, 116(1), 55-79. Retrieved from http://www.jstor.org/stable/2696443

Dana, J., R. A. Weber, et al. (2007). Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. *Economic Theory* 33(1): pp. 67-80.

Ditto, P. H., D. A. Pizarro, et al. (2009). Motivated Moral Reasoning. *Psychology of Learning and Motivation*. D. M. Bartels, C. W. Bauman, L. J. Skitka and D. L. Medin. Burlington, Academic Press. 50: pp 307-338.

Dunning, D., J. A. Meyerowitz, et al. (1989). Ambiguity and self-evaluation: the role of idiosyncratic trait definitions in self-serving assessments of ability. *Journal of Personality and Social Psychology* 57(6): pp. 1082-1090.

Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics*, 643-669.

Anscombe, F. J. & R. J. Aumann. (1963) *The Annals of Mathematical Statistics*, Vol. 34, No. 1. (Mar. 1963), pp. 199-205.

Feiler, L., (2014.) Testing models of information avoidance with binary choice dictator games, *Journal of Economic Psychology*, Elsevier, vol. 45(C), pages 253-267.

Felson, R. B. (1981). Ambiguity and Bias in the Self-Concept. *Social Psychology Quarterly* 44(1): 64-69.

Fudenberg, D., & D. Levine, (2006). A Dual-Self Model of Impulse Control. *The American Economic Review*, 96(5), 1449-1476. Retrieved from http://www.jstor.org/stable/30034982

Golman, R., D. Hagmann & G. Loewenstein. (2017). Information Avoidance. *Journal of Economic Literature*, 55(1): 96-135.

Haisley, E. C. & R. A. Weber (2008). Self-serving interpretations of ambiguity in other-regarding behavior. Department of Social and Decision Sciences, Paper 105: 30 p.

Hetts, J. J., Boninger, D. S., Armor, D. A., Gleicher, F., & Nathanson, A. (2000). The influence of anticipated counterfactual regret on behavior. *Psychology & Marketing*, 17(4), 345-368.

Humphrey, Steven J. (2004) "Feedback-conditional regret theory and testing regret-aversion in risky choice," *Journal of Economic Psychology*, 25(6), 839-857,

Larrick, R. P., & Boles, T. L. (1995). Avoiding regret in decisions with feedback: A negotiation example. *Organizational Behavior and Human Decision Processes,* 63(1), 87-97.

Loomes, G., & Sugden, R. (1982). Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty. *The Economic Journal*, 92(368), 805-824. doi:10.2307/2232669

Lord, C., L. Ross, et al. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology* 37(11): pp. 2089-2109.

Mele, A. R. (1997). Real self-deception. *Behavioral and Brain Sciences* 20(1997): pp. 91-136.

Michel, C. & A. Newen (2010). Self-deception as pseudo-rational regulation of belief. *Consciousness and Cognition* 19(3): 731-744.

Mijović-Prelec, D., & Prelec, D. (2010). Self-deception as self-signalling: a model and experimental evidence. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1538), 227-240.

Quattrone, G. A. & A. Tversky (1984). Causal versus diagnostic contingencies: On self-deception and on the voter's illusion. *Journal of Personality and Social Psychology* 46(2): pp. 237-248.

Ritov, I. (1996). Probability of regret: Anticipation of uncertainty resolution in choice. *Organizational Behavior and Human Decision Processes*, 66(2), 228-236.

Sloman, S. A. & P. M. Fernbach, (2010). Self-deception requires vagueness. *Cognition* 115(2): 268-281.

Szabados, B. (1973). Wishful thinking and self-deception. *Analysis*. 33: 201-205.

Valdesolo, P. & D. DeSteno (2008). The duality of virtue: Deconstructing the moral hypocrite. *Journal of Experimental Social Psychology* 44(5): 1334-1338.

von Hippel, W. & R. Trivers (2011). The evolution and psychology of self-deception. *Behavioral and Brain Sciences* 34: pp 1-56.

Yildiz, M. (2007). Wishful Thinking in Strategic Environments. *The Review of Economic Studies*, 74(1), 319-344. Retrieved from http://www.jstor.org/stable/4123246

Zeelenberg, M., & Pieters, R. (2007). A theory of regret regulation 1.0. *Journal of Consumer Psychology,* 17(1), 3-18.

Zeelenberg, M., J. Beattie, J. van der Pligt, & N. K. de Vries, “Consequences of Regret Aversion: Effects of Expected Feedback on Risky Decision Making,” *Organizational Behavior and Human Decision Processes*, 65(2), p 148-158,