



**GOVERNANCE AND THE EFFICIENCY
OF ECONOMIC SYSTEMS
GESY**

Discussion Paper No. 71

Regret in Dynamic Decision Problems

Daniel Krähmer*
Rebecca Stone**

July 2005

*Daniel Krähmer, Freie Universität Berlin, Institut für Wirtschaftstheorie, Boltzmannstr. 20, 14195 Berlin, Germany,
+49-(0)30-83855223, kraehmer@wiwiss.fu-berlin.de

**Rebecca Stone, ELSE - Department of Economics, University College London, Gower Street, London WC1E 6BT, UK,
+44-(0)20-7679 5894, rebecca.stone@ucl.ac.uk

Financial support from the Deutsche Forschungsgemeinschaft through SFB/TR 15 is gratefully acknowledged.

Sonderforschungsbereich/Transregio 15 · www.gesy.uni-mannheim.de
Universität Mannheim · Freie Universität Berlin · Humboldt-Universität zu Berlin · Ludwig-Maximilians-Universität München
Rheinische Friedrich-Wilhelms-Universität Bonn · Zentrum für Europäische Wirtschaftsforschung Mannheim

Speaker: Prof. Konrad Stahl, Ph.D. · Department of Economics · University of Mannheim · D-68131 Mannheim,
Phone: +49(0621)1812786 · Fax: +49(0621)1812785

Regret in Dynamic Decision Problems*

Daniel Krähmer[†] and Rebecca Stone[‡]

July 22, 2005

Abstract

The paper proposes a framework to extend regret theory to dynamic contexts. The key idea is to conceive of a dynamic decision problem with regret as an intra-personal game in which the agent forms conjectures about the behaviour of the various counterfactual selves that he could have been. We derive behavioural implications in situations in which payoffs are correlated across either time or contingencies. In the first case, regret might lead to excess conservatism or a tendency to make up for missed opportunities. In the second case, behaviour is shaped by the agent's self-conception. We relate our results to empirical evidence.

Keywords: Regret, Counterfactual Reasoning, Reference Dependence, Information Aversion

JEL Classification: C72, D11, D81

*We would like to thank Helmut Bester, Paul Heidhues, Roland Strausz, and especially Erik Eyster for helpful discussions and suggestions. We also thank seminar participants in Berlin, Munich, Vienna and participants of the 2005 ESRC Game Theory Workshop at UCL and the 2005 Toulouse Conference on Economics and Psychology for useful comments. Krähmer gratefully acknowledges support by the German Science Foundation (DFG) through SFB/TR 15.

[†]Freie Universität Berlin, Institut für Wirtschaftstheorie, Boltzmannstr. 20, 14195 Berlin, Germany, ++49-(0)30-83855223, kraehmer@wiwiss.fu-berlin.de.

[‡]ELSE - Department of Economics, University College London, Gower Street, London WC1E 6BT, UK, ++44-(0)20-7679 5894, rebecca.stone@ucl.ac.uk

1 Introduction

A large body of evidence in psychology, economics, neuroscience, and consumer research indicates that the satisfaction people derive from their experiences is influenced by their perception of what would have occurred had they made different choices¹: we experience regret when we perceive that we would have been better off, and we rejoice when we perceive that we are better off than we would have been. Crucially, these emotions are responses to our *ex post* evaluation of our actions relative to unchosen alternatives. Even if our action was optimal given the information available *ex ante*, we nevertheless experience regret when we later discover that some other action would have made us better off. This irrational impulse appears, however, to be coupled with considerable foresight. The evidence suggests that people anticipate these emotions and that they are prepared to trade-off purely material considerations against their desire to avoid future regret (and maximise future rejoicing).

In the light of this evidence, we study the impact of anticipated regret on decision making. We consider an agent who cares directly about the relative performance of his chosen option with respect to foregone alternatives, and who rationally takes these concerns into account when he makes his decisions. This fundamental behavioural assumption is the centre piece of the seminal regret theories of Loomes and Sugden (1982 and 1987a) and Bell (1982 and 1983), which generate predictions that are consistent with many empirical violations of expected utility theory, such as the Allais paradox.²

These theories are essentially static in nature: there is a single decision period, and regrets are realised at the end of the period when the agent learns the state of the world and receives his material payoff. This paper goes beyond these classic treatments and extends regret theory to dynamic decision problems in which multiple decisions have to be taken at different points in time.³

In the analysis of such dynamic situations, a number of conceptually novel issues arise that are absent from a static context. This is because the relative performance of an early decision cannot be fully assessed until the consequences of subsequent decisions that affect the eventual payoff have become apparent. Therefore, the agent's final regrets will only be determined at the culmination of a sequence of actions. This means that later decisions are not only driven by forward looking considerations; they are also influenced by backward looking considerations

¹We review this evidence below.

²A number of studies have experimentally tested the predictions of the classic regret theories. We review this literature below.

³The original formulation of regret theory might be an appropriate tool for analysing decision making in some dynamic contexts. For instance, Irons and Hepburn (2003) use the original formulation of regret theory to analyse a sequential search model.

that arise from the agent's desire to minimise his ultimate regrets from early choices. In other words, in dynamic contexts the regret associated with an action taken in a particular period is generally not "sunk" at the end of that period.

This implies that an agent's regret concerns inevitably lead him to care about his behaviour at counterfactual decision nodes, that is, decision nodes that he would have reached had he acted differently in the past. Consider, for example, a career choice problem in which an agent first chooses whether to study economics or philosophy and then chooses an occupation. Suppose that the agent has chosen economics and is deciding on a career for himself. In doing so, he needs to assess how his current career choice will influence his future regrets that will arise from the comparison between the ultimate payoff consequences of his actual economics degree and his forgone philosophy degree. Now notice that the ultimate payoff that is associated with the decision to study philosophy depends upon the career path that the agent would have subsequently chosen. Hence, the agent's anticipated regret depends upon his conjectures about this counterfactual career path. Moreover, since he cares about regret, these conjectures will influence his optimal career choice as an economist.

We assume that the rationality of the agent imposes constraints on the conjectures he can adopt about this alternative career path. That is, we suppose that he understands that had he chosen philosophy, he would, in order to determine his occupational choice, conjecture in a similar manner about his behaviour as an economist. Thus, we view the agent's period 2 behaviour as the outcome of a game played between the possible selves (economist and philosopher) that arise as a result of his period 1 choice. Our behavioural prediction is that his period 2 behaviour will constitute an equilibrium of this game: a *regret equilibrium*.⁴

Beyond clarifying these conceptual issues, we aim to derive specific behavioural predictions that arise from regret equilibrium. To do so, we require a plausible specification of the agent's preferences. Therefore, before we undertake the analysis of multi-period decision making, we outline a representation of the agent's preferences and demonstrate that, in a static context, it replicates a number of empirical regularities that are suggested by the psychological evidence.

This representation of preferences has two key features. First, the agent cares about the departure of chosen outcomes from a reference point, which is given by his best estimate of what he could have gotten under another choice, given his current information. Second, he dislikes losses relative to the reference point more than he likes equivalent gains. Thus, his fear of regret is a more potent force than his love of rejoicing. This means that the agent is averse

⁴Formally, our game is a special case of a psychological game, and regret equilibrium is a psychological Nash equilibrium (see Geanakoplos et al., 1989). In a psychological game, players care about their beliefs about what other players do and believe. In our game, the agent cares about his beliefs about what his other possible selves do.

to *evaluation risk*: he dislikes uncertainty about the difference between his eventual payoff and reference point.

The implications of this are threefold. First, the agent is averse to payoff uncertainty, holding the reference point constant. Second, he dislikes variation in the reference point, holding the payoff constant. Since new information about unchosen options puts the reference point at risk, this means that he dislikes such information.⁵ Finally, he is better off when the outcome of his chosen action is more positively correlated with the outcome of unchosen alternatives.

In a simple two period model, the agent's regret concerns generate a rich set of behavioural predictions. These results are driven by the fact that the agent may use his second period action to strategically minimise the overall evaluation risk to which he is exposed.

In a setting in which the agent makes two similar decisions in successive periods, we identify two opposing forces on his behaviour. On the one hand, his aversion to feedback on foregone options leads him to behave *conservatively*. He exhibits a tendency to stick to his past choices, even in the face of evidence that indicates that switching would be payoff maximising. This tendency arises because by sticking to his first period choice he avoids learning what he could have received had he taken a different course of action in the first place. On the other hand, he also exhibits a *reparative* tendency, which manifests itself in a tendency to try to make up for missed opportunities. He increases the correlation between his eventual payoff and counterfactual payoff by choosing an action in the second period that is similar to the action that he turned down in the first.

The balance of the two effects depends on how important the second decision is in payoff terms. When it is unimportant, the second period choice has little effect on the extent to which the agent's total payoff matches his counterfactual payoff while its information content is undiminished. This means that his feedback aversion prevails and he acts conservatively. As the second period decision becomes more important in payoff terms, the matching potential of the second period action increases and so he becomes more prone to switching to a different course of action.

In a setting in which the payoff consequences of the first and second period decisions are independent, the agent's prevailing concern in the second period is to match the actions that he believes his counterfactual selves would have chosen in period 2. Suppose, for example, that he has chosen to become a lawyer rather than a consultant and now must decide how hard he will work at his chosen career. Then the belief that he would have worked hard had he become a consultant drives him to work hard as a lawyer, since the pain of doing badly in material terms as a lawyer is then compounded by its unfavourable comparison with the counterfactual

⁵This is in line with psychological evidence, which we review in Section 3.2.

in which he became a successful consultant. Conversely, if he believes that he would have been a lazy consultant he has less incentive to work hard as a lawyer. Thus, the agent’s motivation is influenced by his self-conception, specifically, his beliefs about how he would behave in certain counterfactual situations.

Thus, the agent in our model has counterfactual thought processes which extend beyond the mere enumeration of unchosen options to the consideration of his own behaviour in alternative possible worlds. We go beyond the original regret theories by providing a conceptual framework within which an agent’s counterfactual beliefs and actual behaviour are simultaneously determined as part of an equilibrium of an intrapersonal game.

Related Literature An alternative means of analysing regret concerns in a dynamic context is proposed by Eyster (2002). Eyster derives behavioural implications of a “taste for consistency” whereby agents have a preference for actions which cast past decisions in a more favourable light. Specifically, agents like actions which warrant having taken a past action, which means that agents’ current decisions tend to be influenced by sunk costs (or past opportunity costs). On a related note, Prendergast and Stole (1996) consider a signalling story in which agents overreact to early private information and underreact to later information in order to send a signal to their future selves (or some third party) about the quality of their decisions. Finally, Yariv (2005) considers a model in which agents can actually choose their beliefs directly in order to make themselves feel better about past decisions. Current actions are then taken in the light of the resulting beliefs and so tend to reflect past judgements too closely.

All three of these models share with ours the notion that an agent’s current decisions have a backward looking component. However, these are models in which an agent, in one way or another, explicitly seeks to rationalise past decisions.⁶ In our model, by contrast, the agent has a more basic concern, which is simply to *avoid information* which threatens to cast a past decision in a more unfavourable light. Thus, for instance, switching course of action might reveal information about unchosen alternatives and hence put the reference point at risk. Our agent does not engage in attempts to rationalise past decisions: neither does he have an intrinsic taste for consistency (as in Eyster), nor limited recall that would enable him to influence his future self’s beliefs through costly signalling (as in Prendergast and Stole), nor can he directly choose his beliefs to minimise his regrets (as in Yariv).

From a methodological point of view, the work that is most relevant to the current approach

⁶In Eyster, current actions provide agents with a rationale for actions taken in the past; in Prendergast and Stole, actions rationalise past actions indirectly, insofar as they influence the agent’s beliefs about the wisdom of his past choices; finally, in Yariv’s model, agents directly manipulate their perception of their past decisions.

is Koszegi’s (2004) development of the concept of “personal equilibrium”. Agents are assumed to care directly about their beliefs about the actions they will choose in the continuation games following their current action, and these beliefs are, in turn, required to correctly predict their subsequent equilibrium behaviour.⁷ In contrast, in our setting, the agent cares directly about the beliefs about the actions he would choose in the continuation games following counterfactual actions. Hence, our notion of regret equilibrium may be viewed as an analogue to the notion of a personal equilibrium in situations in which agents care directly about their beliefs about their counterfactual behaviour.⁸

We conclude with a brief overview of the empirical evidence on regret that we alluded to at the outset. First, there is a literature in experimental economics that tested a number of predictions implied by the original regret theories by Loomes and Sugden (1982 and 1987a), such as juxtaposition effects⁹ and violations of monotonicity and transitivity. Early experiments found evidence for these effects and argued that they are best explained by regret theory.¹⁰ Subsequent studies, however, indicate that the results of these early experiments are sensitive to subtle details in the format, in which decision problems are presented to subjects (see Harless, 1992, and Starmer and Sugden 1993, 1998). Thus, the experimental validity of regret theory, as with most choice theories, does appear to be vulnerable to framing effects¹¹: when decision problems are presented in formats, in which regret considerations are made particularly salient, patterns of choice that are uniquely predicted by regret theory do occur and are best explained by regret theory but these effects seem to disappear when the presentation format makes regret considerations less salient (see Starmer and Sugden, 1998).

While this experimental economics literature focusses exclusively on testing predictions of regret theory when the outcomes of unchosen lotteries are always revealed to the agent,

⁷Personal equilibrium includes the notion of loss-aversion equilibrium (see Koszegi and Rabin, 2004, and Heidhues and Koszegi, 2004). This captures “disappointment”, the emotional response to falling short of one’s expectations.

⁸Also related to our approach are a number of papers that consider agents who, like the agents in our paper, care directly about certain beliefs that they hold. For instance, Koszegi (2000a and 2000b) considers implications of an agent’s concern with his self image and Caplin and Leahy (2001) consider preferences which have an anticipatory component, reflecting the pleasure or pain an agent derives from his beliefs about his future prospects.

⁹A juxtaposition effect occurs when preferences over lotteries depend on the joint distribution of the lotteries in the choice set and not only on each lottery’s marginal distribution.

¹⁰See, for example, Loomes (1988a, 1988b), Loomes and Sugden (1987a, 1987b), Loomes et al. (1991, 1992), Starmer and Sugden (1989).

¹¹For example, Starmer and Sugden (1993) find that juxtaposition effects largely disappear when an experimental design is used that controls for so-called event-splitting effects. These effects arise from the tendency of subjects to put more weight on an event if it is presented as several sub-events than if it is presented as a single event.

Bell (1983, p.1165) suggested that the “hypothesis that it may matter whether a foregone alternative is resolved or not ... is the predicted phenomenon on which experimentation should be concentrated”. This proposal has inspired psychologists to conduct experiments to examine the effect that varying the amount of feedback on unchosen alternatives has on subject’s choices. Their findings support a version of regret theory, like the one we use here, in which an agent’s regret concerns make him averse to such feedback.¹²

There is also psychological literature which seeks to measure the impact of the counterfactual emotions of regret and disappointment directly. In a series of experiments Mellers et al. (1999) simultaneously measure emotions - both actual and anticipated - and choice. They measure anticipated emotions by asking people to rate how they would feel about hypothetical monetary outcomes of gambles chosen from a menu of two choices in both the presence and absence of feedback on unchosen gambles. Likewise, actual emotions are measured by asking people to assess their post-choice level of satisfaction when the outcomes are real. Overall, they find support for their hypothesis that agents make choices to maximise their expected anticipated pleasure, where the anticipated pleasure of an outcome depends on its monetary payoff, the extent to which the monetary payoff falls short or exceeds its initially expected payoff (disappointment) and the expected (or actually realised) payoff of the other gambles in the set from which agents may choose (regret).¹³

In addition, there is a literature in consumer research which suggests both that regret influences consumers’ post-choice valuations of their purchases (Inman et al., 1997 and Taylor, 1997) and that anticipation of regret influences purchase decisions (Simonson, 1992).

Finally, there is a literature in neuroscience which seeks to investigate the neurological basis of emotions such as regret and disappointment. For example, using a similar experimental setup to Mellers et al. (1999), Camille et al. (2004) find that, in contrast to normal subjects, the emotional responses of patients with orbitofrontal cortical lesions are insensitive to the nature of the feedback they receive on unchosen gambles, which suggests that they are unable to experience regret. This suggests that the experience of regret is closely associated with a particular region of the brain (the orbitofrontal cortex).¹⁴

The rest of the paper is organised as follows. In Section 2, we outline the basic model of preferences that captures the agent’s regret concerns. In Section 3, we derive the essential properties of these preferences that drive the main results of the paper and consider some basic implications for behaviour when there is just a single decision period. In Section 4, we extend

¹²We review this evidence in more detail in Section 3.2.

¹³See also Mellers (2000) and Mellers and McGraw (2001).

¹⁴Disappointment effects were mitigated somewhat, but were nevertheless present, suggesting that regret has a different neurological basis from that of disappointment.

the model to two decision periods and introduce the concept of regret equilibrium. In Section 5, we apply the equilibrium concept to the analysis of behaviour in two different scenarios: (1) a setting in which the agent makes similar decisions in each of the two periods; and (2) a setting in which payoffs from the second decision are independent of the payoffs from the first. Section 6 concludes. Throughout, we relate our findings to the psychological evidence. All proofs are relegated to an appendix.

2 The static setup

In this section, we outline the basic model that we use to analyse decision making in a single period. There is a single agent who chooses an action from a finite set Y . An action $y \in Y$ corresponds to a real-valued random variable, a *lottery*, X_y that is distributed with c.d.f F_y . A typical realisation of X_y is denoted x_y , and the set of all lotteries is denoted by $L = \{X_y | y \in Y\}$. We assume that the lotteries in L have finite first moments. Further assumptions are mentioned in passing.

The agent cares about two things. First, he cares about his material payoff. When the agent chooses lottery X_y , then his instantaneous material payoff from realisation x_y is given by $\phi(x_y)$, where ϕ is a real-valued, continuous, increasing and weakly concave function. Thus, ϕ has the properties of a standard utility function in which marginal material utility is declining in wealth.

Second, he cares about his evaluation of the performance of his action relative to the performance of forgone alternatives in his choice set. When he perceives that his action has performed relatively well, he rejoices and his utility improves. When he perceives that it has performed relatively poorly, he regrets his decision and his utility falls.

The agent's evaluation of an action's relative performance is given by the deviation of its material utility from a reference point, R , that represents the agent's aggregate assessment of the performance of other actions in Y . To capture this formally, we divide the timeline into a decision period, $t = 1$, during which the agent chooses an action y , and an evaluation period, $t = T > 1$. In the evaluation period, the agent receives his material payoff and evaluates the relative performance of y given his information in $t = T$.

To simplify the exposition, we assume in this section that in period T the agent is either fully informed about an option in Y or not informed about it at all.¹⁵ This means that the agent's period T information can be captured by a subset $L_T \subseteq L$ of lotteries whose realisations are known to the agent in period T . More specifically, we assume that L_T contains a set ΔL_T

¹⁵We shall relax this assumption when we consider the dynamic context. All results for the single period problem extend easily to the case where the agent's information in T is generated by general signals.

of lotteries whose realisations are revealed to the agent after he has chosen y in period 1 as well as the outcome of the lottery chosen in $t = 1$. Hence, $L_T = \Delta L_T \cup \{X_y\}$. We shall write $I(L')$ to denote the information set generated by the subset $L' \subseteq L$ and denote by $I_T = I(L_T)$ the agent's information in period T .¹⁶

The reference point against which the agent evaluates his decision to choose y is then given by the agent's estimate of the material utility he could have obtained from the best possible unchosen alternative in his choice set given his final information. Thus,

$$R_y = \max_{y' \in Y \setminus y} E[\phi(X_{y'}) | I_T].$$

Notice that R_y is a real-valued random variable whose realisation the agent learns in period T at the latest.¹⁷ We denote a typical realisation of R_y by r_y .¹⁸

In the evaluation period, given a reference point realisation r_y , the agent evaluates the difference $\phi(x_y) - r_y$ between his actual material utility and his reference point by a real-valued, continuous, increasing, and strictly concave function ρ with $\rho(0) = 0$, whose domain ranges over all possible values of $\phi(x_y) - r_y$.¹⁹ That is, the agent's overall instantaneous utility from action y , which is realised in period T , is defined as

$$u(y) = \phi(x_y) + \theta \rho(\phi(x_y) - r_y),$$

where the parameter $\theta \geq 0$ measures the intensity of the agent's regret concerns relative to his material concerns. We call an agent with $\theta > 0$ a *regretful* agent and an agent with $\theta = 0$ a *standard* agent.

$\rho(0) = 0$ reflects the fact that the agent likes a positive relative performance (he rejoices) and dislikes a negative relative performance (he experiences regret). The concavity of ρ implies *loss aversion*, that is the agent's regret about a negative performance is stronger than his rejoicing about an equivalently positive performance.²⁰

Three comments about this specification are in order. First, it should be clear that the regretful agent evaluates his decision from an ex-post perspective in light of his final information.

¹⁶More precisely, I_T is the σ -algebra or partition - defined on an appropriate state space - that is generated by the lotteries in L_T . The expectation conditional on I_T is then the expectation conditional on knowing the realisations of all lotteries in L_T .

¹⁷Formally, R_y also depends on Y . Since we consider cases only in which Y is fixed, we suppress this dependency.

¹⁸We have chosen the maximum specification mainly for ease of exposition. An alternative specification is a weighted average over the expected values of the forgone options. All of our results carry over to this case.

¹⁹Strict concavity means that for all $\lambda \in (0, 1)$, x, x' : $\rho(\lambda x + (1 - \lambda)x') > \lambda \rho(x) + (1 - \lambda)\rho(x')$.

²⁰Loss aversion is, by now, a familiar and well-established empirical regularity (Rabin, 1998). In the context of regret, e.g. Mellers et al. (1999), Camille et al. (2004), and Inman et al. (1997) find that regret looms larger than rejoicing.

This is the fundamental behavioural assumption of regret theory. The agent regrets a decision with a poor outcome even though this decision might have been optimal at the point of decision making. Indeed, if the agent evaluated the relative performance of his decision based upon the information that was available to him at the point of decision making, then his behaviour would be indistinguishable from that of a standard expected utility maximiser.

Second, our framework builds on Bell (1983) who, like us, allows for the possibility that unchosen lotteries are only partially resolved. He looks at independent lotteries and provides natural conditions under which the reference point is given by the expectation of the unchosen lottery. In contrast, Loomes and Sugden (1982) and Bell (1982) implicitly assume that the uncertainty with respect to forgone lotteries is fully resolved after the choice was made. In this case, the reference point is determined by the realisation of the unchosen lotteries, and our formulation reproduces the original formulations.

Third, concavity is a familiar assumption in economics, but less so in psychology. Indeed, Mellers et al. (1999) suggest that the regret function is S -shaped, that is ρ is concave over gains (rejoicing domain) but convex over losses (regret domain).²¹ We assume concavity of ρ because it is an analytically elegant way to generate a behavioural regularity that features prominently in a number of experiments on regret: people tend to make choices that minimise their exposure to information on forgone alternatives.²²

3 Decision making in the static setup

This section has two purposes. First, we outline basic properties of the regretful agent's preferences when there is a single decision period in order to develop intuitions that are helpful in understanding the dynamic decision problems studied below. Second, we argue for the plausibility of our specification of preferences by deriving behavioural predictions and showing that they are in line with empirical evidence.

Throughout we assume that the agent correctly anticipates his emotional response at the evaluation stage and maximises his expected utility taking his regret concerns into account. This means that the agent chooses an action $y \in Y$ so as to maximise his expected instantaneous utility

$$U(y, \Delta L_T) \equiv E[u(y)] = E[\phi(X_y)] + \theta E[\rho(\phi(X_y) - R_y)].$$

This specification of preferences has the crucial implication that the regretful agent is averse to *evaluation risk*: the ex ante variability of the difference between the payoff of his chosen

²¹This is reminiscent of the S-shaped value function of Kahneman and Tversky's (1979) Prospect Theory.

²²Section 3.2 will make these points explicit.

action and his reference payoff. This is the key property of the agent’s preferences that we shall explore in this paper. Evaluation risk aversion is a direct consequence of the concavity of the regret function ρ . The following Lemma is immediate (and thus stated without proof).

Lemma 1 *Let $y, y' \in Y$ and let R_y and $R_{y'}$ be the corresponding reference points. Suppose $E[\phi(X_y)] = E[\phi(X_{y'})]$ and $E[R_y] = E[R_{y'}]$, and let $\phi(X_y) - R_y$ second order stochastically dominate $\phi(X_{y'}) - R_{y'}$. Then the agent prefers y to y' .*

Evaluation risk aversion implies that the agent is averse both to ex ante variation in his payoff holding his reference point constant, and to variation in his reference point holding his payoff constant. We refer to the former effect as *payoff risk aversion* and the latter effect as *reference point risk aversion*. In addition, it implies that the agent likes the lotteries in his choice set to be positively correlated. We refer to this effect as the *correlation effect*.

In the following subsections, we formally characterise these effects and consider what they imply about a regretful agent’s overall risk attitudes.

3.1 Payoff risk aversion

In the absence of feedback on foregone options, the agent behaves like an ordinary risk-averted. This is summarised formally by the following Lemma.

Lemma 2 *Suppose that $(X_y)_{y \in Y}$ are stochastically independent. Suppose there is no exogenous feedback, i.e. $\Delta L_T = \emptyset$. Fix two actions, $y, y' \in Y$, with $E[\phi(X_y)] = E[\phi(X_{y'})]$. Let X_y second order stochastically dominate $X_{y'}$. Then the agent prefers y to y' . This is true even if ϕ is linear.*

Lemma 2 is a straightforward consequence of the concavity of ρ and ϕ . It is consistent with psychological evidence which suggests that anticipated regret may be associated with more risk averse decision making.²³

From a methodological point of view, Lemma 2 highlights the importance of distinguishing clearly between behaviour that is motivated by regret and behaviour that is motivated by standard considerations of diminishing marginal utility. Therefore, we must isolate behavioural

²³Zeelenberg (1996) reports that “implicit or explicit in most experimental work on anticipated regret is the idea that it leads to risk aversion” (p. 149). According to Kardes (1994), “concern about regret that may follow a bad decision promotes extreme risk aversion” (p. 448). In a consumer research context, Simonson (1992) finds that experimental subjects make more risk averse choices (paying a higher price for a better known brand) when asked to anticipate regret and responsibility. Finally, Josephs et al. (1992) provide experimental evidence that low self-esteem subjects who, they argue, may be more predisposed to experience regret tend to make more risk averse choices than high self-esteem agents, who may be less predisposed to such feelings.

predictions that arise in our framework that could not be generated by expected utility theory alone. Hence, we take the behaviour of the standard agent as the benchmark against which we contrast the regretful agent's behaviour.

3.2 Reference point risk aversion

Evaluation risk aversion also implies reference point risk aversion. The following Lemma characterises what we mean by this.

Lemma 3 *Fix an action $y \in Y$ and consider two final information sets I_T and I'_T . Suppose that I'_T is more informative than I_T . Then the agent prefers final information set I_T to final information set I'_T .*

Intuitively, the agent dislikes uncertainty about his future reference point for any given action he has chosen because such uncertainty increases the fluctuations in the relative performance of his chosen action. Since reference point uncertainty arises only when the agent receives information about unchosen options after he has made his choice, it follows that the agent is averse to receiving such information. In this sense, reference point risk aversion corresponds to *ex post information aversion*.

One of the central findings of psychological research is that regret concerns induce risk preferences that depend upon the feedback on unchosen options that people expect to receive.²⁴ The following example illustrates that reference point risk aversion generates this feature.

Example 1 Let $Y \in \{a, b\}$. Suppose that $\phi(x) = x$ and that X_a and X_b are both symmetrically distributed with $E[X_a] = E[X_b]$. Suppose that the agent learns the realisation of X_a regardless of his choice while he learns the realisation of X_b only if he chooses b . Then the agent (weakly) prefers a to b .

Example 1 says that, all else equal, the agent chooses actions that minimise his exposure to ex-post information. This is straightforward consequence of reference point risk aversion: by choosing a , the agent's reference point is safe whereas by choosing b , he faces ex ante uncertainty about his reference point. Notice that the result holds irrespective of the risk associated with each option. Even if X_b second order stochastically dominates X_a , the regretful agent chooses a . Moreover, it follows from Lemma 2 that the agent would choose the less risky option

²⁴On this point, see also Bell (1983) who considers the choice between a risky gamble and a sure thing and hypothesises that the sure thing becomes more attractive if the uncertainty about the risky gamble is not resolved (p. 1160).

if no feedback was provided on unchosen options, and it is easy to see that the same is true if he always receives feedback on all options.

The relation between risk preferences and feedback has been tested experimentally in a number of psychological studies. In support of our model, most experimental studies find that subjects' behaviour is influenced by the amount of feedback they expect to obtain on unchosen options.²⁵ Larrick and Boles (1994), in a negotiation experiment, and Ritov (1996), in a gambling experiment, find that subjects become more willing to take risks if feedback on the risky alternative is provided. Similar findings are reported by Zeelenberg et al. (1996) in a gambling task, and by Zeelenberg and Beattie (1997) in an ultimatum game experiment.²⁶ For a more detailed review of the experimental evidence see Zeelenberg (1999).

In a field study, Zeelenberg and Pieters (2004) find that anticipated regret is a predictor of participation among players of the Dutch Postcode Lottery, but not among players of the State Lottery. The difference is attributed to the fact that winners of the Postcode Lottery are determined by postcode, and thus, potential participants learn whether or not they would have won regardless of whether they choose to play from their neighbours, while such forced feedback about the outcome is not a feature of the State Lottery.

3.3 A correlation effect

In the previous subsections, we derived implications of evaluation risk aversion in situations in which either the agent's payoff or his reference point was held constant. We now look at situations in which payoff and reference point are interdependent. This is the case when lotteries are correlated. In order to provide a parsimonious means of varying the degree of correlation between lotteries, we compare a simple lottery with a compound lottery. More precisely, for two lotteries X and \widehat{X} , and for $\beta \in [0, 1]$ we denote by $\beta \circ X \oplus (1 - \beta) \circ \widehat{X}$ the compound lottery that yields lottery X with probability β and lottery \widehat{X} with probability $(1 - \beta)$. When $\beta = 1$, the compound lottery is perfectly positively correlated with X , and when $\beta = 0$, the compound lottery is perfectly correlated with \widehat{X} .

As the following Lemma makes clear, the agent's regret concerns cause him to like the lotteries in his choice set to be positively correlated. We refer to this effect as the *correlation effect*.

Lemma 4 *Let X_a and \widehat{X}_a be i.i.d.²⁷. Fix $\beta \in [0, 1]$, and define $X_\beta = \beta \circ X_a \oplus (1 - \beta) \circ (\widehat{X}_a)$.*

²⁵An exception is Kelsey and Schepanski (1991) who, in the context of a taxpayer reporting decision experiment, find no evidence that the feedback that subjects expect to obtain has an impact on behaviour.

²⁶In a gambling experiment, Josephs et al (1992) induce risk averse choices of low self-esteem subjects by providing feedback on forgone options, but in their setup, the risk averse choice minimises regret.

²⁷identically and independently distributed.

Let $Y = \{a, \beta\}$ and suppose the agent chooses a . Then the agent's utility increases in β . This is true regardless of whether or not the outcome of X_β is revealed.

The intuition for Lemma 4 is straightforward. Evaluation risk aversion means that the agent is averse to ex ante fluctuations in the relative performance of his chosen action. Positive correlation between his actual payoff and his reference payoff reduces these fluctuations. In the extreme case of perfect positive correlation, any variation between an action's payoff and its reference payoff is eliminated, and the agent is perfectly insured against any regrets.²⁸

4 Dynamic decision making and regret equilibrium

We now turn to the analysis of dynamic decision problems in which the ultimate payoff the agent receives is the consequence of multiple decisions. In analysing such decision problems, two key issues arise. First, an agent's subsequent decision may influence the regret he experiences about early actions. Second, his conjectures about his behaviour at counterfactual decision nodes play a crucial role in determining his eventual regrets, and, hence, his actions.




Figure 1: Two decision trees

To illustrate the first issue, consider the following decision problem that may be represented by decision tree I in Figure 1. In period 1, the agent has to invest money in either a risky project (R) or a safe opportunity (S). In period 1.5, if he has chosen the risky project, he obtains a signal indicating whether the project is likely to pay off or not, and in period 2 he has to decide whether to continue the project (c) or to abandon it (a). If he continues the project, he

²⁸It is straightforward to establish an analogous result, which shows that the agent's utility is decreasing in the degree of negative correlation between two lotteries in his choice set.

obtains the return from the risky project (which might be either better or worse than the safe opportunity). If he abandons it, he loses his initial investment (but avoids losses if the project does not pay off). All payoffs are realised in period 3.

Suppose that the agent has chosen R , received an unfavourable signal about the risky project and now has to decide between c and a . On the one hand, given his current information, R appears to have been the wrong initial choice and thus the agent might already regret having chosen R . On the other hand, the comparison between R and S only becomes fully evident in period 3 when the ultimate payoffs are realised. Indeed, if he continues the project, it might pay off with some likelihood in which case the agent will rejoice in R rather than regret it. Therefore, the regrets associated with R should arise—at least partially—in period 3.

This simple observation has an important implication: in period 2, the regrets that the agent expects to experience in period 3 as a result of choosing R will depend on whether the agent chooses c or a in period 2. If he abandons the project, he will definitely regret having chosen R rather than S . If he chooses to continue, however, he will rejoice in R with some likelihood.

Now, consider a slight modification of the previous example. Suppose that the return from S is also risky and, moreover, correlated with the return from R . Suppose also that, when the project is abandoned after choice of R , no further information about R or S is revealed. In this case, the decision between c and a affects not only the payoff from R but also the reference point against which R is evaluated. If c is chosen, information about S will be revealed (due to correlation between R and S) and the agent will update his reference point, while if a is chosen, this will not be the case.

In general, when the agent decides between c and a in period 2, the eventual regrets from his R decision are not “sunk” because they can still be affected by his second period decision. In other words, the second period decision will have a backward looking component that arises from the agent’s desire to minimise his ultimate regrets from his first period choice. To capture this, we shall assume that regrets about a particular decision are realised at an ultimate evaluation stage, after which no further events which may affect the apparent ex post wisdom of the decision will occur. Making this assumption is the natural way to extend classic regret theory to dynamic decision problems: as we previously noted, classic regret theory is built around the assumption that agents evaluate their past decisions in the light of any information that they have at their disposal *once the payoffs from their choices and any available information about unchosen alternatives has been realised*.

To illustrate the role that the agent’s counterfactual conjectures will play, consider the set of decision problems that may be represented by tree II. Such decision problems have the new feature that the final payoff of *both* initial choices depends upon a continuation action. Suppose

that the agent has already chosen A and is contemplating his second period decision. In doing so, he needs to determine how his second period decision affects the extent to which he will regret the initial choice of A in period 3. The extent to which he will regret this choice will depend upon what he perceives that he would have gotten had he chosen B instead of A . This depends on the continuation action that the agent believes he would have chosen after B . So the agent's regrets about having chosen A will depend upon his conjecture about what he would have done in the counterfactual situation that would have arisen had he chosen B in the first period. And given such a conjecture, the performance of A relative to that of B will depend upon whether he chooses a_A or b_A .

Thus, to complete the specification of dynamic decision problems with regret, we need to specify how the agent forms his conjectures about his behaviour at counterfactual decision nodes. In the next section, we develop a framework that resolves this problem in the spirit of rational expectations. For simplicity, we restrict the analysis to two-stage problems. The extension to an arbitrary finite number of stages is conceptually similar.

4.1 The dynamic setup

We extend the timeline to two decision periods, $t = 1, 2$, and one evaluation period $t = T > 2$. In period 1, the agent chooses an action $y_1 \in Y_1$. After having chosen y_1 in period 1, the agent is exogenously informed about the realisation z_1 of a random variable (*signal*) $Z_1(y_1)$. Denote by $I_2 = I(Z_1(y_1))$ the agent's information set when he makes his period 2 choice.²⁹

In period 2, contingent on having chosen y_1 in period 1 and on signal realisation z_1 , the agent chooses action $y_2 \in Y_2(y_1)$. The set of feasible actions is denoted by $Y = \{(y_1, y_2) | y_1 \in Y_1, y_2 \in Y_2(y_1)\}$, and each action $y \in Y$ has an associated lottery X_y . The set of all lotteries is denoted by $L = \{X_y | y \in Y\}$. After having chosen y_2 the agent is exogenously informed about the realisation z_2 of a signal $Z_2(y_1, y_2)$. In the evaluation period, the agent receives his material payoff and evaluates his decision based on the information available in $t = T$, which is given by $I_T = I(Z_1(y_1), Z_2(y_1, y_2), X_{(y_1, y_2)})$.³⁰

This specification contains the simplifying assumption that *all* regrets are realised in the evaluation period and rules out the existence of interim regrets that might be experienced at stage 2. As we argued above, the important point of the dynamic setup is that some regrets from the period 1 decision are realised in an ultimate evaluation period. Our assumption focusses our attention exclusively on the effect that these ultimate regret concerns have on the agent's

²⁹As before, we write $I(Z)$ to denote the information set generated by the signal Z .

³⁰The specification of I_T assumes that the agent has perfect memory and is informed about the outcome of his choice in the final period.

decision making.

4.2 Preferences

When an agent makes decisions in two different periods, he might regret (or rejoice in) each decision. Thus, for example, an agent may regret his initial choice of degree while also rejoicing in his subsequent career choice, believing that his chosen career was the best option available to him given his (albeit imperfect) choice of degree. To capture this, we disentangle the agent's overall regret from action (y_1, y_2) into two components, one corresponding to each decision. Accordingly, his final payoff is evaluated with respect to two different reference points, R_1 and R_2 .

In period T , when the agent evaluates his second decision, y_2 , he takes his period 1 action, y_1 , as given. That is, he compares the performance of y_2 with the performance of the alternative actions $y'_2 \in Y_2(y_1)$ that were available to him at the period 2 decision node induced by action y_1 . Hence, his second period reference point, R_2 , is determined in exactly the same way as it would have been had he simply been required to choose from the same choice set, $Y_2(y_1)$, in a single period context. That is, the instantaneous regret that the agent experiences about his second decision y_2 in the evaluation period is given by

$$v_2(y_1, y_2) = \rho \left(\phi(x_{(y_1, y_2)}) - \max_{y'_2 \in Y_2(y_1) \setminus y_2} E \left[\phi \left(X_{(y_1, y'_2)} \right) \middle| I_T \right] \right)$$

As we noted previously, the key feature of the dynamic setup is that the outcome of a forgone action $y'_1 \neq y_1$ depends on which period 2 action the agent would choose in the counterfactual scenario that arises in the event that he has chosen y'_1 in period 1. Thus, to compare the performance of his actual choice y_1 to that of a counterfactual choice y'_1 , the agent needs to form a conjecture about his counterfactual continuation action at the decision node induced by y'_1 . We identify action y'_1 with a *counterfactual self* $S(y'_1)$: the incarnation of the agent that would have controlled the period 2 action in the event that y'_1 had been chosen in period 1.

Given that the agent chose y_1 , the agent's conjecture is expressed by a *belief*, $\mu_2^{(y'_1, z_1)}$, that specifies for each counterfactual self, $S(y'_1)$, $y'_1 \neq y_1$, and all possible realisations z_1 of $Z_1(y'_1)$ a probability distribution over $Y_2(y_1)$. That is, the agent conjectures that in the event that counterfactual self $S(y'_1)$ observes signal $z_1 = z_1(y'_1)$, it chooses action $y_2 \in Y_2(y'_1)$ with probability $\mu_2^{(y'_1, z_1)}(y_2)$. We denote the profile of all (contingent) beliefs about counterfactual strategies by $\mu_2^{-y_1} = \left(\mu_2^{y'_1} \right)_{y'_1 \in Y_1 \setminus y_1}$.

The relative performance of action y_1 with respect to a forgone alternative y'_1 is determined as follows. The agent first computes the conditional probability that the counterfactual self $S(y'_1)$ is of type z_1 conditional on his ex-post information I_T . Using this he derives the probability

$\mu_2^{(y'_1, z_1)}(y_2)$ with which $S(y'_1)$ has chosen action y_2 and then he uses these probabilities to average over the expected values of $\phi\left(X_{(y'_1, y_2)}\right)$ conditional on I_T .

The agent's first period reference point, R_1 , is then given by the performance of the best possible unchosen alternative in his first period choice set. Formally, given a conjecture $\mu_2^{-y_1}$ and given actions (y_1, y_2) , the instantaneous regret that the agent experiences about his first decision y_1 in the evaluation period is given by

$$v_1(y_1, y_2; \mu_2^{-y_1}) = \rho \left(\phi(x_{(y_1, y_2)}) - \max_{y'_1 \in Y_1 \setminus y_1} E_{z_1(y'_1)} \left[E_{y_2} \left[\phi\left(X_{(y'_1, y_2)}\right) \middle| I_T \right] \middle| I_T \right] \right),$$

where the first expectation is taken with respect to the conditional distribution of the signal $Z_1(y'_1)$ conditional on I_T , and the second expectation is taken with respect to the agent's conjecture $\mu_2^{(y'_1, z_1(y'_1))}$ and the conditional distribution of $\phi\left(X_{(y'_1, y_2)}\right)$ conditional on I_T .

Finally, we make the assumption that the total regret that the agent experiences is a weighted sum of his regret about his first and second decision. That is, the total instantaneous regret that the agent experiences in the evaluation period as a result of actions (y_1, y_2) is given as

$$v(y_1, y_2; \mu_2^{-y_1}) = \theta_1 v_1(y_1, y_2; \mu_2^{-y_1}) + \theta_2 v_2(y_1, y_2),$$

where the parameters $\theta_t \geq 0$ measure the extent to which the agent cares about his first and second period decision in regret terms.

Notice that our formulation embodies the psychological assumption that if the agent believes that his counterfactual self chooses his second period action randomly, either because it plays a mixed strategy, or because it makes its choices contingent on the realisation of a private signal, he incorporates this uncertainty directly into the reference point. Since the agent is averse to information about the reference point, he prefers not to receive information about what exactly the counterfactual self ends up choosing. For instance, he dislikes information about the realisation of the private signal that the counterfactual self has received if the counterfactual self plays a contingent strategy. This assumption is in line with our assumptions about the agent's informational preferences in the static setup.

4.3 Behaviour

We now describe the assumptions that we make about the agent's behaviour. We argue that in this dynamic setup the agent's behaviour is naturally described as the outcome of a non-cooperative game played between the possible incarnations of the agent (actual and counterfactual) that arise at each of his period 2 decision nodes.³¹

³¹Throughout we assume that the agent cannot commit to a period 2 action contingent on his first period choice. If this was the case, we would be back in the static case: the agent would simply select a contingent

Consider first the agent's period 2 behaviour and suppose that his period 1 choice was y_1 . We make two assumptions. First, we assume that, *given* a conjecture $\mu_2^{-y_1}$, the agent maximises his expected utility taking into account his regret concerns. That is, he chooses $y_2 \in Y_2(y_1)$ so as to maximise

$$U_{y_1}(y_2; \mu_2^{-y_1}) \equiv E[\phi(X_{(y_1, y_2)}) + v(y_1, y_2; \mu_2^{-y_1}) | I_2].$$

Implicit in this formulation is the assumption that the agent is bound to his conjecture over time. That is, in the evaluation period, the agent's conjecture about his counterfactual self's strategy is the same as his period 2 conjecture.

Second, we assume that in forming his conjectures he rationally understands that had he chosen y'_1 , he would, in order to determine his behaviour in this counterfactual situation, reason in a similar way about his current behaviour. This suggests that the agent's period 2 behaviour is the outcome of an intra-personal game G played between the agent's counterfactual selves where counterfactual self $CS(y_1)$'s objective function is given by U_{y_1} .

Because each self cares directly about his conjectures about his counterfactual selves' actions, the game G is formally recognisable as a psychological game (see Geanakoplos et al., 1989) and we are naturally led to assume that the agent's period 2 behaviour constitutes a psychological Nash equilibrium. More precisely, let $s_2^{y_1} \in Y_2(y_1)$ be a strategy for self $S(y_1)$, and denote by $s_2 = (s_2^{y_1})_{y_1 \in Y_1}$ a strategy profile and by $\mu_2 = (\mu_2^{y_1})_{y_1 \in Y_1}$ a profile of conjectures. We define a regret equilibrium as a psychological Nash equilibrium of G .

Definition 1 *A pair (μ_2, s_2) is a (period 2) regret equilibrium if*

(i) $\mu_2 = s_2$

(ii) for each $y_1 \in Y_1$ and $s'_2 \in Y_2(y_1)$: $U_{y_1}(s_2; \mu_2^{-y_1}) \geq U_{y_1}(s'_2; \mu_2^{-y_1})$.

The set of all strategies s_2^ such that there is a (period 2) regret equilibrium (s_2^*, s_2^*) is denoted by S_2^* .*

Our key behavioural assumption is that the agent's period 2 behaviour constitutes a (period 2) regret equilibrium. Thus, we assume that each self plays a best response to his conjecture about the other self's behaviour and these conjectures correctly predict the other selves' behaviour.³² Note that (i) implies that an equilibrium is fully specified by s_2 alone. Thus, henceforth we omit any explicit mention of μ_2 .

strategy in period 1 and compare it to other contingent strategies that he could have chosen in period 1. In the absence of a commitment device, we think this is a natural assumption.

³²Since the two selves operate in different possible worlds, the actions of one do not affect the other's payoffs directly. One self's actions affect the other only because they are correctly anticipated in equilibrium and each self cares directly about their beliefs about the other self's behaviour.

Next consider period 1 behaviour. We assume that in period 1 the agent anticipates that he will play a regret-equilibrium in period 2 and chooses action y_1 so as to maximise

$$U(y_1) \equiv E \left[\phi \left(X_{(y_1, s_2^{y_1})} \right) + v \left(y_1, s_2^{y_1}; s_2^{-y_1} \right) \right]$$

s.t. $s_2 \in S_2^*$.

We call the agent's overall behaviour (y_1, s_2) with an optimal period 1 action y_1 and $s_2 \in S_2^*$ a *regret equilibrium*.

The assumption that the agent's behaviour constitutes a regret equilibrium is, in effect, an assumption about the thought process via which the agent constructs the counterfactuals against which he compares actual outcomes. The equilibrium concept embodies the assumption that the agent cannot choose his conjectures freely, but must form them rationally based on his understanding of the nature of the choices that his alternative possible selves must make. That is, he views his counterfactual selves as fully rational agents who, just like himself, are motivated by comparisons of their own actions with those of similarly rational counterfactual selves. In short, he forms his conjectures in a *sophisticated* manner.

Alternatively, we could have assumed that the agent naively supposes that his counterfactual selves simply choose whichever action maximises their expected material utility. More radically, we could have assumed that the agent chooses his beliefs in a self-serving manner so as to minimise his regrets.³³ We do not consider the effects of such alternative assumptions here, since we want to focus on behavioural implications of anticipated regret alone. Keeping assumptions about the rationality of the agent in tact isolates the effects of this single departure from the standard model. Also, the assumption of sophistication here seems to be the natural extension of the notion of rational expectations to a setting in which the agent cares directly about his beliefs about counterfactual outcomes.³⁴

³³Given that the agent's conjectures are the outcome of a purely mental process about events which never, in fact, happen, it is natural to wonder whether the agent might engage in a kind of "wishful thinking" and choose his beliefs about his counterfactual selves in order to make himself feel good about his actual choices. Such an agent's beliefs would be irrational in a truth-seeking sense, but not necessarily irrational from the point of view of maximising the agent's overall utility.

³⁴One could plausibly argue that the appropriate extension of rational expectations would restrict agents conjectures about their opponents to rationalisable strategies rather than equilibrium strategies. In particular, in the current context a player can never observe his opponents' actions ex-post and thus a justification of equilibrium play by a learning argument fails. In this sense, our equilibrium notion is too strong. However, such an objection can also be levelled more generally against ordinary economic models in which rational agents are assumed to anticipate equilibrium play (on this point, see Bernheim, 1982, pp. 1008-1010). Moreover, as Koszegi (2004) argues, the common knowledge requirements necessary for equilibrium play might be more easily satisfied in an intra-personal game when different selves of the same individual interact than in an ordinary game (p. 8).

5 Implications of regret equilibrium

We now derive behavioural implications of our regret equilibrium concept. Notice that the strategic interaction between the agent's possible selves arises through the agent's regret concerns with respect to his period 1 choices. We therefore set $\theta_2 = 0$ from now on in order to isolate these strategic effects. For $\theta_2 > 0$, all our results go through so long as period 1 regret concerns are sufficiently large relative to period 2 regret concerns. Also, we are mainly interested in the agent's behaviour at stage 2. We do not explicitly look at his—rather straightforward—first period behaviour.

We consider two kinds of situation. First, we consider repeated actions whose payoffs are correlated across time. Second, we look at situations in which the payoffs of period 2 actions are correlated across all possible decision nodes.

5.1 Correlation across time

Suppose the agent can choose repeatedly from actions with correlated payoffs. We identify two opposing forces that drive the agent's incentives in period 2. On the one hand, the agent's aversion to information on forgone lotteries will drive him to stick to past actions, that is, to behave in an *excessively conservative* manner. On the other hand, his desire to correlate his actual outcome with his counterfactual reference outcome will drive him to try to make up for opportunities that he turned down in early periods by switching his course of action, a tendency that might psychologically be interpreted as a *reparative tendency*.

In the next two subsections, we identify conditions under which each effect is likely to occur and summarise the relevant evidence. We consider the following setup.³⁵ In each of two periods, $t = 1, 2$, the agent has to make a binary choice $y_t \in Y_t = \{a, b\}$. The payoffs of an action are perfectly correlated across time. That is, action a yields lottery X_a in period 1 and lottery λX_a in period 2 for $\lambda \geq 0$. Likewise, action b yields lottery X_b in period 1 and lottery λX_b in period 2. Lotteries X_a and X_b are stochastically independent. After the first period, the agent receives a signal $Z = (Z_a, Z_b)$ that is informative about X_a and X_b .

The parameter λ reflects the weight on the second period payoff relative to the first period payoff. When analysing the agent's behaviour, we will be interested in the effects of varying λ .

5.1.1 Excess conservatism

We now consider how a regretful agent responds to news that questions the wisdom of his first period choice. We assume that once the agent has made his second period choice, he only learns

³⁵The setup can be easily generalised without changing our main results.

the payoffs of the actions that he has chosen, but receives no additional information about the other action's payoffs. Hence, the agent's terminal information set is $I_T = I(Z, X_{y_1}, X_{y_2})$. This assumption is crucial, because it enables the agent to manipulate the amount of information he ultimately receives as a result of his second period decision.

More specifically, we assume that X_a and X_b have zero mean and compact support. We consider signals that reveal whether an action's payoff is positive or negative. Define the signal Z_{y_1} as

$$Z_{y_1} = \begin{cases} +1 & \text{if } X_{y_1} \geq 0 \\ -1 & \text{if } X_{y_1} < 0. \end{cases}$$

After receiving the signal, the agent updates his beliefs about an action's payoff. Conditional on observing signal realisation z_{y_1} the agent's posterior about X_{y_1} is the resulting truncated distribution. We assume that this conditional distribution is non-degenerate for all signal realisations (z_a, z_b) .³⁶ That is, the signal is not perfectly revealing.

Suppose the agent has chosen a in period 1 and received signal realisation (z_a, z_b) . Then we shall say that news is "bad" if $z_a < z_b$. If he has chosen b in period 1, a similar definition applies.

The agent's second period behaviour is described by a strategy $s_{y_1} = s_{y_1}(z_a, z_b)$ that prescribes an action in Y_2 contingent on all signal realisations. We are interested in how the agent reacts to news. We shall say that the agent's behaviour is *conservative* if he does not change his action in response to "bad" news, i.e. if $s_a(-1, +1) = a$ and $s_b(+1, -1) = b$.

Before we start the analysis of the agent's second period behaviour, we briefly consider the benchmark case of a standard agent.

Lemma 5 *The period 2 behaviour of a standard agent is never conservative.*

The Lemma holds more generally for signals that alter the agent's interim beliefs about the two options. In general, a standard agent seeks a compromise between increasing his expected payoff and lowering the aggregate risk to which he is exposed. Intuitively, switching to a different action in the second period reduces the aggregate risk, because it diversifies the agent's portfolio. Hence, the agent sticks to his first period choice only if it has an advantage in terms of expected return. But this can never be the case if he receives bad news.

We now characterise the regretful agent's second period behaviour and identify conditions such that it exhibits conservatism.

Proposition 1 *There exist strictly positive $\bar{\lambda}$ and $\bar{\theta}_1$ such that for all $\lambda \leq \bar{\lambda}$ and $\theta_1 \geq \bar{\theta}_1$ there is a regret-equilibrium in which the agent's period 2 behaviour is conservative.*

³⁶This is the case if X_a and X_b each contain more than one positive and more than one negative point in their support.

The key point of Proposition 1 is that in period 2 the regretful agent may stick to his first period choice even though he receives unfavourable news and switching his action would both raise his material payoff and reduce the aggregate payoff risk to which he is exposed. In this sense, the agent acts as if he ignores evidence against his previous choice and exhibits *excess conservatism*.

To understand what drives this result, notice that for a very small period 2 payoff weight, λ , the period 2 decision has negligible payoff consequences while its information content is undiminished. Thus, because the signal Z is not perfectly informative, it amounts to rejecting or accepting ex-post information about the alternative forgone in period 1. Switching course of action reveals what the agent would have obtained had he acted differently in period 1 and exposes the agent to reference point risk whereas staying with his first period choice adds nothing to his final information set and so eliminates all reference point risk.

This conservative tendency provides an explanation for brand loyalty: by continuing to consume a particular brand, consumers avoid the risk of learning that they would have been better off had they consumed a different brand. In effect, the fear of regret generates a psychological switching cost. At a general level, as Camerer and Weber (1999) demonstrate, there is evidence that people and organisations sometimes exhibit a tendency to maintain or increase a commitment to a project even when the cost-benefit calculus suggests that they would be wise to abandon it. Our theory suggests one factor which could contribute to such escalation phenomena.

In fact, Tykocinski and Pittman (1998) offer some experimental confirmation of the idea that people tend to stick to an original course of action in order to avoid facing the fact that they ought to have chosen a different course of action in the first place. Specifically, the results of their experiments suggest that when an individual initially forgoes an attractive action opportunity, he will tend to decline a similar but substantially less attractive current opportunity, even though it still has positive value.

Moreover, they argue that such behaviour is driven by a fear of regret.³⁷ These experiments do not, however, exactly replicate the features of our setup since subjects perfectly learn the value of the option initially forgone before making their subsequent choice. However, in explaining their results, Tykocinski and Pittman implicitly assume that agents are able to forget such information about unchosen options that they may have once known unless they are explicitly reminded of it. Thus, agents stick to an original course of action in order to avoid *reminding* themselves of the value of the forgone alternative. We assume, by contrast, that agents have

³⁷“Declining the subsequent action opportunity may appear to be illogical because this opportunity is still attractive in an absolute sense, but it may also be psychologically rewarding because it mitigates the unpleasant experience of regret produced by dwelling on a perceived loss” (p. 608).

perfect memory but that they remain imperfectly informed about an option in the event that they did not choose it. Thus, agents stick to an original course of action in order to avoid *learning* about a forgone option.

Interestingly, Tycocinski and Pittman also find that this tendency to avoid the foregone course of action is attenuated when the agent will be reminded of it even if he does not subsequently choose it. Similarly, in our setup, if we make the alternative assumption that the regretful agent will receive full feedback on his available actions, irrespective of the choices that he makes, then he no longer displays conservatism.

5.1.2 Reparative action

The previous result rests critically on the assumption that the second decision's payoff consequences are rather unimportant. In this case, the regretful agent's desire to avoid information about foregone alternatives drives him to stick to his first period action in the second period. In this subsection, we identify a countervailing force on his behaviour, which drives him to try to make up for missed opportunities by choosing a second period action that is similar to the action he turned down in the first period. This force predominates when the two decisions are of equal importance in payoff terms.

In contrast to the previous section we now assume, for simplicity, that the agent does *not* get any feedback after period 1. The behavioural tendencies we identify will, however, also be generated by a more complex model in which the agent does receive such interim feedback.

It turns out that the relative riskiness of the options a and b is a critical determinant of the agent's incentives. Thus, in order to characterise the forces at work, we shall want to vary the degree of risk of a relative to b . A convenient way to do so, is to assume that X_a is symmetrically distributed with mean $E[X_a]$, and that X_b is given by the random variable $\gamma\widehat{X}_a$, where \widehat{X}_a and X_a are i.i.d. The parameter γ reflects the relative riskiness of a and b . If $\gamma \geq 1$, then X_b second order stochastically dominates X_a , and the reverse if $\gamma < 1$.

In addition, our point becomes most transparent if $\lambda = 1$. So, for the rest of this section, we focus exclusively on this case. We also assume that ϕ is strictly concave.³⁸

Our aim is to identify conditions under which the behaviour of the regretful agent departs from that of the standard agent. We thus first characterise the standard agent's behaviour in the current setup.

Lemma 6 *Consider the standard agent. There are $\overline{\gamma}, \overline{\overline{\gamma}}$ with $\overline{\gamma} < 1 < \overline{\overline{\gamma}}$ such that the agent's*

³⁸The case with linear ϕ is a boundary case, which is qualitatively similar, but requires slight changes in the proofs. To save space, we omit the argument.

choice is given by

$$(y_1, y_2) = \begin{cases} (b, b) & \text{if } \gamma < \bar{\gamma} \\ (a, b) \text{ or } (b, a) & \text{if } \gamma \in [\bar{\gamma}, \bar{\bar{\gamma}}) \\ (a, a) & \text{if } \gamma \geq \bar{\bar{\gamma}}. \end{cases}$$

The lemma says that the standard agent switches course of action if and only if the lotteries are similarly risky, that is, if $\gamma \in [\bar{\gamma}, \bar{\bar{\gamma}})$. (Due to the symmetry of our setup, he is indifferent between (a, b) and (b, a)). In this case, by switching course of action the agent diversifies the aggregate payoff risk to which he is exposed. As the relative riskiness of one lottery becomes extreme ($\gamma < \bar{\gamma}$ or $\gamma \geq \bar{\bar{\gamma}}$), then choosing the less risky lottery in both periods minimises the agent's overall payoff risk.

We now turn to the regretful agent. The following lemma describes the best responses given the counterfactual self switches action in period 2.

Lemma 7 *Consider the regretful agent's actual period 2 decision.*

- (i) *Suppose the agent chose a in period 1 and the counterfactual self switches action in period 2. Then there is $\hat{\gamma}$ with $\hat{\gamma} > \bar{\bar{\gamma}}$ such that the agent switches course if and only if $\gamma < \hat{\gamma}$.*
- (ii) *Suppose the agent chose b in period 1 and the counterfactual self switches action in period 2. Then there is $\hat{\gamma}$ with $\hat{\gamma} < \bar{\gamma}$ such that the agent switches course if and only if $\gamma > \hat{\gamma}$.*

The following proposition describes the equilibria that result from these best responses. (It is an immediate consequence of the previous lemma and thus stated without proof.)

Proposition 2 *For all $\gamma \in (\hat{\gamma}, \hat{\bar{\gamma}})$ there is a regret equilibrium in which the regretful agent switches action in period 2.*

The proposition implies that the regretful agent switches for a greater range of parameters than the standard agent. To grasp the intuition for this result, notice first that the regretful agent's payoff risk aversion means that, like the standard agent, he has an incentive to diversify his portfolio. However, there is an additional force at work on the regretful agent's behaviour. Recall from the static analysis that the regretful agent dislikes fluctuations between his actual payoff and payoffs that would have resulted from alternative courses of action. If the weight on the action's payoff is the same across periods and the counterfactual strategy prescribes switching in period 2, then if the agent switches in period 2, he obtains in period 1 what he counterfactually would have obtained in period 1. Hence, by switching he eliminates any fluctuations between his actual payoff and his corresponding reference point.

One interpretation of the regretful agent's behaviour is suggested by psychological "dissonance" theories that view the mere act of choice as inherently painful due to the regret that is

caused by the necessity of having to give up certain opportunities (see Festinger, 1964). The agent in our model makes up for the opportunity missed in the first period, and thereby eliminates the regrets that necessarily arise from being forced to choose an action in period 1. In this sense he appears to exhibit a “reparative” tendency as he seeks to relieve the dissonance associated with his period 1 choice.

We close this section with an informal discussion of what happens if the weight on an action’s payoff is different across periods ($\lambda \neq 1$). In this case, the agent’s opportunities to obtain this type of insurance against regret are limited. On the one hand, if λ is small, then the second period action cannot contribute much to reduce the fluctuations between overall actual and counterfactual payoffs. As in the previous section, the period 2 incentive is then determined by the agent’s feedback concerns and there is a regret equilibrium in which the agent sticks to his first period action.

On the other hand, if λ is large, the main fluctuations in overall relative performance come from period 2 payoffs. Thus, the highest correlation between actual and counterfactual payoff can be achieved by matching one’s counterfactual period 2 action. Thus, while a standard agent optimally chooses the second period action whose associated lottery is less risky, a regretful agent may optimally choose an action whose associated lottery is stochastically dominated if he believes that his counterfactual self also chooses this action. It follows from that when λ is large, asymmetric equilibria can arise in which the agent chooses the same action in period 2 irrespective of this first period action. This matching incentive is at the core of the results of the next section, to which we now turn.

5.2 Correlation across counterfactuals

In this section, we consider a scenario in which the payoffs from first and second period actions are independent and the payoffs from second period actions are correlated across all possible decision nodes. We first derive a result that identifies a tendency of the regretful agent to match the actions that he believes his counterfactual self would choose in period 2. As an application of this general tendency we then consider a career choice example and show that the same agent might make very different equilibrium effort choices in period 2, depending on whether he holds optimistic or pessimistic beliefs about his counterfactual behaviour.

The setup is as follows. In period 1, the agent’s choice set is $Y_1 = \{a, b\}$ with associated lotteries X_a and X_b . In period 2, his choice set is $Y_2^{y_1} = \{c, d\}$.³⁹ We denote the lottery associated with a period 2 choice y_2 by $X_{y_2}^{y_1}$, where the superscript y_1 indicates that the lottery depends on which action y_1 was chosen in period 1. We assume that $X_{y_2}^a$ and $X_{y_2}^b$ have the

³⁹Strictly speaking, c and d depend on y_1 . We suppress this dependency for notational simplicity.

same distribution and are stochastically independent of X_a and X_b . The key point is that $X_{y_2}^a$ and $X_{y_2}^b$ might be correlated. To capture this, let X_{y_2} and \widehat{X}_{y_2} be i.i.d. and let $\alpha \in [0, 1]$. Then we set

$$X_{y_2}^a = X_{y_2}, \quad X_{y_2}^b = \alpha \circ X_{y_2} \oplus (1 - \alpha) \circ \widehat{X}_{y_2}.$$

That is, if $\alpha = 1$, the payoffs from actions are perfectly correlated across decision nodes, and if $\alpha = 0$, they are stochastically independent. The agent's overall material payoff from path (y_1, y_2) is given by $X_{(y_1, y_2)} = X_{y_1} + X_{y_2}^{y_1}$. Finally, we assume for simplicity that the agent does not receive interim information and that the outcome of unchosen lotteries is always revealed to the agent.⁴⁰

5.2.1 Matching the counterfactual self

We now describe the agent's period 2 incentives. To isolate the effects that arise as result of the correlation in period 2, we focus on the case, in which the first stage is inconsequential in terms of material payoff. That is, we assume that X_a and X_b are deterministic and equal to zero. Notice that this does *not* mean that the first decision is inconsequential in terms of regret. Rather, it is precisely the distinctive feature of our approach that the regrets about the first decision also include the regrets about the second period action opportunities that arise as a consequence of this first decision.

We analyse the interplay between the relative riskiness of the two options c and d and the degree of correlation across decision nodes. To do so, we assume that X_d second order stochastically dominates X_c and vary the correlation parameter α .

Notice that the standard agent chooses d at the second stage for all α . However, for the regretful agent this does not need to be true. Our first result says that if lotteries are sufficiently correlated across decision nodes, the agent has an incentive to match his counterfactual behaviour.

Lemma 8 *Consider the agent's actual period 2 decision and suppose his counterfactual self chooses action y_2 in period 2. Then there is a $\widehat{\theta}_1 > 0$ such that for all $\theta_1 \geq \widehat{\theta}_1$, there is an $\widehat{\alpha} \in [0, 1]$ such that for all $\alpha \geq \widehat{\alpha}$, it is a best response to choose y_2 , too.*

To illustrate the intuition for this result, consider the case, in which the counterfactual self chooses the riskier option c . The actual self then faces a trade-off between risk-minimisation and regret-minimisation. In terms of risk, option d is superior to option c by assumption. In terms of regret, option c is superior to option d . This is once more due to the agent's desire to minimise the fluctuations in his overall relative performance. Because payoffs are correlated

⁴⁰In this section, the nature of the feedback the agent receives is unsubstantial.

across decision nodes, he can reduce these fluctuations by choosing c , that is, by aligning his second period actions across decision nodes. When the counterfactual self chooses d , then the same two forces both favour option d for the actual self.

Proposition 3 is a direct consequence of the agent’s best response (and thus stated without proof).

Proposition 3 *Let θ_1 and α be sufficiently large. Then there are two period 2 regret equilibria, one in which the agent always chooses c in period 2, and one in which he always chooses d in period 2 irrespective of his period 1 choices.*⁴¹

Thus, the agent has an incentive to act in a way that he would have behaved under different circumstances. We explore implications of this point further in the following career example.

5.2.2 Application: belief determined motivation

Suppose that in period 1, the agent chooses between two occupations, a and b , where, for concreteness, a corresponds to “consultancy” and b corresponds to “the law”. In the second stage, the agent chooses an effort level $e \in \{0, 1\}$ at cost ke , $k > 0$, that influences the likelihood of his career success. The key point is that we think of success to be determined by some general unknown ability of the agent that is the same in both occupations. That is, we suppose that the likelihood of success is correlated across occupations. Formally, the agent’s ability is given by a state $\omega \in \{0, 1\}$ where each state’s prior probability is $1/2$. In state ω , the agent’s material payoff (gross of effort costs) is ωe in both occupations. For simplicity, we assume that the agent learns his true ability after his effort choice.

The matching incentive gives rise to a complementarity between the agent’s conjecture about his counterfactual effort choice and his own effort choice. To see this formally, define Δ_e as the agent’s period 2 incentive to choose high effort (the difference between the utility he obtains from working hard and the utility he obtains from shirking) given that his counterfactual self chooses effort e . Then we have:

Proposition 4 *The agent’s incentive to choose high effort is higher the higher the counterfactual self’s effort, i.e. $\Delta_1 \geq \Delta_0$.*

The result is a straightforward consequence of the matching incentive. It shows that the agent’s motivation depends on what he conjectures he would have done in his counterfactual occupation. If as a consultant the agent believes that he would have worked hard as a lawyer, he better work hard as a consultant so as not to regret not having become a successful lawyer. In other

⁴¹It is easy to see that, generically, these two equilibria are the only pure-strategy equilibria.

words, the agent’s motivation is determined by his general views about himself: his self-image. An agent who is generally optimistic about himself with respect to other possible worlds, will tend to be more motivated than a more pessimistic but otherwise identical agent.⁴²

The idea that motivation is driven by self-perceptions is familiar to psychologists. In their influential review of the evidence, Taylor and Brown (1988) conclude that “self-enhancing perceptions ... appear to foster motivation, persistence at tasks, and ultimately, more effective performance.” (p. 199) and argue that differences in self-images across individuals can be explained as the result of a sort of self-fulfilling prophecy: someone who is optimistic (pessimistic) about his talents will tend to be more (less) motivated, resulting in better (worse) outcomes. Better (worse) outcomes, in turn, will confirm his optimistic (pessimistic) expectations. For an elaboration of this view, see Aspinwall et al. (2001).

Our approach shares with these accounts the theme of self-fulfilling expectations. However, in our model, it is not expectations about actual future outcomes but about counterfactual comparisons that determine current incentives. We are not aware of psychological studies that have pointed to the motivating force of this type of counterfactual reasoning.⁴³

6 Conclusion

In this paper, we have developed an extension of regret theory, which can be used to analyse dynamic decision problems. The key idea is to conceive of a dynamic decision making problem with regret as an intra-personal game in which the agent forms rational conjectures about the behaviour of the various counterfactual selves that he could have been.

Our approach raises a series of issues that go beyond the analysis that we have undertaken in this paper. For example, it is easy to see that in a static setup, a regretful agent might be better off with a smaller choice set because it reduces the number of alternatives that he has to forgo. In our dynamic setup we can ask whether the agent would actually choose to pre-commit to less choice given that, after the fact, he might regret his first period commitment.

⁴²Of course, in equilibrium, effort choice and self-image have to be consistent. It is straightforward to see that multiple equilibria might arise in our setup, one in which the agent works hard in both occupations and one in which he chooses zero effort in both occupations. Thus, even though the agent must form his beliefs about his counterfactual selves in a rational way in our setup, it is compatible with different, self-sustaining self-images and behaviours.

⁴³Markus and Nurius (1986) introduce the notion of possible selves and argue that the desire or fear to become a specific future self provides incentives for current action. While they do emphasise the importance of counterfactual comparisons for construing the *current* self-image (p. 963), they maintain that the attractiveness of a future self is determined by its actual features only but not by the counterfactual comparisons it will engage in.

At a more conceptual level, our theory highlights the need to investigate further the close interaction between regret and counterfactual reasoning. In classic regret theory, the agent engages in comparisons between what he actually did and what he *possibly could* have done counterfactually. We have argued, however, that once we move into the realm of dynamic decision problems, the agent’s reference point is inevitably determined by his conjectures about what he *reasonably would* have done at a second stage had he made a different first period choice.

Psychological research suggests that such considerations might play a role not only in dynamic but also in static decision problems. Indeed, Kahneman and Miller (1986) argue that counterfactual comparisons depend in important ways on people’s ability to mentally “undo” an antecedent event that led to the actual outcome. For example, regrets arising from poor outcomes are typically larger if a modification of the antecedent event is easier to imagine. In this spirit it seems reasonable to assume that people’s regrets about their own choices should depend on how easy it is for them to perceive that they would have acted differently in the past. Regrets might then arise as a result of the discrepancy between the agent’s actual choice and of what he thinks he would have chosen could he choose again on the basis of his ex-post knowledge. We plan to explore these conceptual issues in a companion paper.

Finally, evidence suggests that, in reality, the experience of regret interacts in important ways with the experience of disappointment that is generated by the discovery that reality has fallen short of expectations (Roese, 1997). Combining our model with the previously mentioned models of personal equilibrium is likely to be a fruitful way to understand this interaction in more detail.

Appendix

Proof of Lemma 2 Suppose that X_y second order stochastically dominates $X_{y'}$. We need to show that the agent’s preference for y over y' , defined as $\Delta U_y = U(y, \Delta L_T) - U(y', \Delta L_T)$, is non-negative. ΔU_y can be written as

$$E[\phi(X_y) - \phi(X_{y'})] + \theta E[\rho(\phi(X_y) - R_y) - \rho(\phi(X_{y'}) - R_{y'})], \quad (1)$$

where, since $\Delta L_T = \emptyset$, $R_y = \max_{\hat{y} \in Y \setminus y} E[\phi(X_{\hat{y}}) | X_y]$ and $R_{y'} = \max_{\hat{y} \in Y \setminus y'} E[\phi(X_{\hat{y}}) | X_{y'}]$. Now observe first that since ϕ is weakly concave, and X_y second order stochastically dominates $X_{y'}$, the first term of (1) is non-negative. It remains to show that the second term is also non-negative. To see this, notice that because $(X_y)_{y \in Y}$ are stochastically independent, the reference points R_y and $R_{y'}$ are in fact constants, and because $E[\phi(X_{y'})] = E[\phi(X_y)]$, it follows that

$R_y = R_{y'}$. Hence, we are done if we can show that

$$\theta E [\rho(\phi(X_y) - R_y) - \rho(\phi(X_{y'}) - R_y)] \quad (2)$$

is non-negative. Define an auxiliary function, $g(x) = \rho(\phi(x) - R_y)$. Since ρ and ϕ are concave, g is also concave. Since X_y second order stochastically dominates $X_{y'}$, it follows that (2) is non-negative which is what we sought to prove. \square

Proof of Lemma 3 Consider information sets $I_T = I(\Delta L_T)$ and $I'_T = I(\Delta L'_T)$ where ΔL_T is contained in $\Delta L'_T$. Let R_y and R'_y be the agent's reference point given final information sets I_T and I'_T respectively. We need to show that the preference for information set I_T over I'_T , defined as $\Delta U_I = U(y, \Delta L_T) - U(y, Y, \Delta L'_T)$, is non-negative. ΔU_I is given by

$$\Delta U_I = \theta (E[\rho(\phi(X_y) - R_y)] - E[\rho(\phi(X_y) - R'_y)]).$$

The law of iterated expectations applied on this expression yields that it suffices to show that

$$E[\rho(\phi(X_y) - R_y) - \rho(\phi(X_y) - R'_y) | I_T] \geq 0. \quad (3)$$

To see that (3) holds, notice that since X_y and R_y are fully determined by I_T , it follows that

$$E[\rho(\phi(X_y) - R_y) | I_T] = \rho(\phi(X_y) - R_y).$$

Moreover, since ρ is concave, Jensen's inequality implies that

$$E[\rho(\phi(X_y) - R'_y) | I_T] \leq \rho(E[\phi(X_y) | I_T] - E[R'_y | I_T]). \quad (4)$$

Hence, (3) follows if we can show that the right hand side of (4) is smaller than $\rho(\phi(X_y) - R_y)$. To see this, recall first that $R'_y = \max_{y' \in Y \setminus y} E[\phi(X_{y'}) | I'_T]$. Hence, since $\max_{y' \in Y \setminus y} \{\phi(X_{y'})\}$ is a convex function in the vector $\{\phi(X_{y'})\}_{y' \in Y \setminus y}$, it follows from Jensen's inequality that $E[R'_y | I_T] \geq \max_{y' \in Y \setminus y} E[E[\phi(X_{y'}) | I'_T] | I_T] = R_y$. Furthermore, $E[\phi(X_y) | I_T]$ equals $\phi(X_y)$ and ρ is increasing. Thus, the right hand side of (4) is smaller than $\rho(\phi(X_y) - R_y)$, which is what we wanted to show. \square ⁴⁴

Proof of Example 1 We have to show that the agent's incentive to choose a , defined as $U(a) - U(b)$, is non-negative. Since $E[X_a] = E[X_b]$ and $\phi(x) = x$, this incentive is given by

$$\theta E[\rho(X_a - E[X_b | X_a])] - \theta E[\rho(X_b - X_a)].$$

⁴⁴This proof is the only point where we make use of the maximum specification of the reference point. If the reference point were given by a linear combination of the expected payoffs of the unchosen alternative, the result would still hold. In this case, the law of iterated expectations implies that $E[R'_y | I_T] = R_y$.

To see that this is non-negative, notice first that since X_a and X_b are symmetrically distributed, the second term in this expression equals $\theta E[\rho(X_a - X_b)]$. Moreover, Jensen's inequality for conditional expectations implies that

$$\rho(X_a - E[X_b|X_a]) \geq E[\rho(X_a - X_b)|X_a].$$

Hence, the first term is not larger than $\theta E[\rho(X_a - X_b)]$. These two observations imply the claim. \square

Proof of Lemma 4 Suppose first that the agent does not receive feedback on X_β . Then, his reference point is $E[\phi(X_\beta)|X_a] = \beta\phi(X_a) + (1 - \beta)E[\phi(\widehat{X}_a)]$. Thus, his utility from choosing a is given by

$$u(a) = E[\phi(X_a)] + E\left[\rho\left((1 - \beta)\left(\phi(X_a) - E\left[\phi(\widehat{X}_a)\right]\right)\right)\right].$$

Now notice that since $E[X_a] = E[\widehat{X}_a]$, the mean of the lottery $(1 - \beta)\left(\phi(X_a) - E\left[\phi(\widehat{X}_a)\right]\right)$ does not change in β , but its risk falls in β . Thus, the concavity of ρ implies that u increases in β .

Suppose next that the agent receives feedback on X_β . Then with probability β , the agent's reference point is $\phi(X_a)$ and with probability $1 - \beta$, his reference point is $\phi(\widehat{X}_a)$. Thus, his utility from choosing a is

$$u(a) = E[\phi(X_a)] + (1 - \beta)E\left[\rho\left(\phi(X_a) - \phi(\widehat{X}_a)\right)\right].$$

By Jensen's inequality and because X_a and \widehat{X}_a are identically distributed: $E\left[\rho\left(\phi(X_a) - \phi(\widehat{X}_a)\right)\right] \leq 0$. Hence, u increases in β . \square

Proof of Lemma 5 To prove the claim, we consider the case only in which the agent chose a in period 1. (Identical considerations apply to the case in which he chose b .) Under this assumption, we compare the utilities from choosing a and b in the second period. Suppose the agent has observed realisation $(Z_a, Z_b) = (-1, +1)$. Then his expected utility from choosing a is given by

$$E[u(a, a)] = E[\phi(X_a + \lambda X_a) | X_a < 0, X_b > 0],$$

and his expected utility from choosing b is

$$E[u(a, b)] = E[\phi(X_a + \lambda X_b) | X_a < 0, X_b > 0].$$

We want to show that the agent's incentive to choose b is positive. But notice that on the event $\{X_a < 0, X_b > 0\}$, the value of $X_a + \lambda X_b$ is strictly larger than that of $X_a + \lambda X_a$. Therefore,

since ϕ is increasing $E[u(a, b)]$ is clearly larger than $E[u(a, a)]$. \square

Proof of Proposition 1 We shall look for conditions such that in period 2, the agent sticks to his first period action for all types of news, that is, $s_a(z_a, z_b) = a$ for all $z_a, z_b \in \{-1, +1\}$ is a best response against $s_b(z_a, z_b) = b$ for all $z_a, z_b \in \{-1, +1\}$, and thus (s_a, s_b) is a period 2 regret-equilibrium. Suppose that a was chosen in period 1 and suppose that $s_b(z_a, z_b) = b$ for all $z_a, z_b \in \{-1, +1\}$. (The case in which b was chosen in period 1 is symmetric.)

If the agent chooses a in the second period, then because X_a and X_b are independent, he learns nothing more about X_b after making his second period decision. Given his conjecture about his counterfactual behaviour, the reference point R_1 against which he evaluates his first decision is the expected payoff from the path (b, b) conditional on the realisation Z . Hence, $R_1 = E[\phi((1 + \lambda) X_b) | Z]$. Using these reference points and the payoff risk associated with a , the expected utility from choosing a in period 2 is given by

$$E[u(a, a)] = E[\phi((1 + \lambda) X_a) | Z] + \theta_1 E[\rho(\phi((1 + \lambda) X_a) - E[\phi((1 + \lambda) X_b) | Z]) | Z].$$

If the agent chooses b in the second period, he will eventually learn both X_a and X_b . Therefore, his reference point is $R_1 = \phi((1 + \lambda) X_b)$, and his expected utility from choosing b is given by

$$E[u(a, b)] = E[\phi(X_a + \lambda X_b) | Z] + \theta_1 E[\rho(\phi(X_a + \lambda X_b) - \phi((1 + \lambda) X_b)) | Z].$$

Thus, the agent's incentive to choose a , $E[u(a, a)] - E[u(a, b)]$, can be written as the sum $\xi_0 + \xi_1$, where the first part is the difference in the material payoff:

$$\xi_0 = E[\phi((1 + \lambda) X_a) | Z] - E[\phi(X_a + \lambda X_b) | Z],$$

and the second part is the difference in regret with respect to the first decision:

$$\begin{aligned} \xi_1 = \theta_1 (& E[\rho(\phi((1 + \lambda) X_a) - E[\phi((1 + \lambda) X_b) | Z]) | Z] \\ & - E[\rho(\phi(X_a + \lambda X_b) - \phi((1 + \lambda) X_b)) | Z]) \end{aligned}$$

We want to show that for any realisation of Z , the sum over the ξ 's is positive if λ is small and θ_1 is large. To do so, notice that ξ_0 is bounded from below. (This is, because the support of X_a and X_b is compact and ϕ is continuous, thus bounded on the support of X_a and X_b .) Hence, it suffices to show that for any realisation of Z , ξ_1 is strictly positive for small λ . We can then choose θ_1 large enough such that ξ_1 outweighs the possibly negative impact of ξ_0 .

To see that ξ_1 is strictly positive for small λ , we first subtract and add

$E[\rho(\phi((1 + \lambda) X_a) - \phi((1 + \lambda) X_b)) | Z]$ to ξ_1/θ_1 to obtain

$$\begin{aligned} \xi_1/\theta_1 = & E[\rho(\phi((1 + \lambda) X_a) - E[\phi((1 + \lambda) X_b) | Z]) - \rho(\phi((1 + \lambda) X_a) - \phi((1 + \lambda) X_b)) | Z] \\ & + E[\rho(\phi((1 + \lambda) X_a) - \phi((1 + \lambda) X_b)) - \rho(\phi(X_a + \lambda X_b) - \phi((1 + \lambda) X_b)) | Z]. \end{aligned}$$

Define the difference in the first line of the right hand side as $\delta(\lambda)$. Jensen's inequality yields that $\delta(\lambda) > 0$. The inequality holds strict, because ρ is strictly concave and the conditional distribution of X_a conditional on Z is non-degenerate. Define $\delta_{\min} = \min_{\lambda \in [0,1]} \delta(\lambda)$. Since the minimum is taken over a compact set, we have that δ_{\min} is strictly positive. We thus obtain:

$$\xi_1/\theta_1 > \delta_{\min} + E[\rho(\phi((1+\lambda)X_a) - \phi((1+\lambda)X_b)) - \rho(\phi(X_a + \lambda X_b) - \phi((1+\lambda)X_b)) | Z].$$

To conclude the proof, it therefore remains to show that the second term in this expression is not smaller than $-\delta_{\min}$ for small λ . To see this, define the arguments in ρ respectively as

$$\begin{aligned} \Delta_1^\phi(\lambda, x_a, x_b) &= \phi((1+\lambda)x_a) - \phi((1+\lambda)x_b), \\ \Delta_2^\phi(\lambda, x_a, x_b) &= \phi(x_a + \lambda x_b) - \phi((1+\lambda)x_b). \end{aligned}$$

Notice first that $\Delta_1^\phi(\lambda, x_a, x_b) - \Delta_2^\phi(\lambda, x_a, x_b) \rightarrow 0$ for $\lambda \rightarrow 0$. Moreover, because the support of X_a and X_b is compact, $\Delta_1^\phi(\lambda, x_a, x_b) - \Delta_2^\phi(\lambda, x_a, x_b)$ is uniformly continuous for all $\lambda \leq 1$. It follows from these two observations that for all $\tau > 0$ there is a $\lambda_\tau > 0$ such that for all realisations x_a, x_b :

$$\left| \Delta_1^\phi(\lambda, x_a, x_b) - \Delta_2^\phi(\lambda, x_a, x_b) \right| < \tau \text{ for all } \lambda \leq \lambda_\tau.$$

In addition, ρ is uniformly continuous on the compact support of X_a and X_b . Hence, there is a $\tau > 0$ such that for all realisations of Δ_1^ϕ and Δ_2^ϕ with $|\Delta_1^\phi - \Delta_2^\phi| < \tau$: $\rho(\Delta_1^\phi) - \rho(\Delta_2^\phi) > -\delta_{\min}/2$. This implies that if we now choose $\lambda \leq \lambda_\tau$, then

$$E[\rho(\phi((1+\lambda)X_a) - \phi((1+\lambda)X_b)) - \rho(\phi(X_a + \lambda X_b) - \phi((1+\lambda)X_b)) | Z] > -\delta_{\min}/2,$$

which is what we sought to prove. \square

Proof of Lemma 6 Suppose first that $\gamma \geq 1$. Notice first that the choice $(y_1, y_2) = (b, b)$ is dominated by the choice (a, a) . This is so, because (b, b) is associated with the lottery $2X_b = 2\gamma\widehat{X}_a$, and the choice (a, a) is associated with the lottery $2X_a$. But since X_a and \widehat{X}_a are i.i.d., $2X_a$ second order stochastically dominates $2\gamma\widehat{X}_a$ for $\gamma \geq 1$.

Moreover, the incentive to choose (a, b) or (b, a) rather than (a, a) is given by

$$\xi_0(\gamma) = E\left[\phi\left(X_a + \gamma\widehat{X}_a\right)\right] - E[\phi(2X_a)]. \quad (5)$$

We shall now show that (a) $\xi_0(\gamma)$ is declining in γ with $\lim_{\gamma \rightarrow \infty} \xi_0(\gamma) = -\infty$, and (b) $\xi_0(1) > 0$. This then implies that there is a $\bar{\gamma} > 1$ such that the agent chooses (a, b) if and only if $\gamma < \bar{\gamma}$, which is the claim that we need to show.

We begin with (b). This follows from the fact that $X_a + \widehat{X}_a$ second order stochastically dominates $2X_a$ and because ϕ is strictly concave.

As for (a), we use the symmetry of X_a and \widehat{X}_a to write

$$\begin{aligned} E \left[\phi \left(X_a + \gamma \widehat{X}_a \right) \right] &= \int_0^\infty \int_0^\infty [\phi(- (x_a + \gamma \widehat{x}_a)) + \phi(x_a + \gamma \widehat{x}_a)] \\ &\quad + [\phi(- (x_a - \gamma \widehat{x}_a)) + \phi(x_a - \gamma \widehat{x}_a)] dF_a(x_a) dF_a(\widehat{x}_a). \end{aligned}$$

The strict concavity of ϕ implies that for all x_a, \widehat{x}_a the terms in the squared brackets are negative and converge monotonically to $-\infty$ as $\gamma \rightarrow \infty$. Thus, the Monotone Convergence Theorem implies that the integral converges monotonically to $-\infty$ as $\gamma \rightarrow \infty$. This establishes the claim for $\gamma \geq 1$. The case $\gamma < 1$ follows with identical arguments. \square

Proof of Lemma 7 To establish part (i), suppose that $s_b = a$, and that the agent has chosen a in period 1. Then, if he chooses a in period 2, he does not learn X_b and his reference point is $R_1 = E[\phi(X_b + X_a)]$. Hence, with $X_b = \gamma \widehat{X}_a$, the expected utility from choosing a in period 2 is

$$u(a) = E[\phi(2X_a)] + \theta_1 E \left[\rho \left(\phi(2X_a) - E \left[\phi \left(X_a + \gamma \widehat{X}_a \right) \right] \right) \right].$$

If he chooses b in period 2, he learns X_b and his reference point is $R_1 = \phi(X_b + X_a)$. Note that his reference point is equal to his actual overall payoff. Thus, the agent will have no regrets and his expected utility is

$$u(b) = E[\phi(2X_a)].$$

Let

$$\xi_1(\gamma) = E \left[\rho \left(\phi(2X_a) - E \left[\phi \left(X_a + \gamma \widehat{X}_a \right) \right] \right) \right],$$

then the incentive to switch to b in period 2 is given by

$$\Delta_b(\gamma) = u(b) - u(a) = \xi_0(\gamma) - \theta_1 \xi_1(\gamma),$$

where $\xi_0(\gamma)$ is defined by (5). We now show that (a) $\Delta_b(\gamma)$ is declining in γ with $\lim_{\gamma \rightarrow \infty} \Delta_b(\gamma) = -\infty$, and (b) $\Delta_b(0) > 0$, and (c) $\Delta_b(\overline{\gamma}) > 0$, where $\overline{\gamma}$ is the cutoff at which the standard agent is indifferent between (a, b) and (a, a) . This then implies that there is a $\widehat{\gamma} > \overline{\gamma}$ such that the regretful agent switches to b if and only if $\gamma < \widehat{\gamma}$, which is the claim that we want to show.

We begin with (b) and (c) and the observation that Jensen's inequality implies that

$$\xi_1(\gamma) < \rho \left(E[\phi(2X_a)] - E \left[\phi \left(X_a + \gamma \widehat{X}_a \right) \right] \right) = \rho(-\xi_0(\gamma)),$$

where the inequality is strict, since ρ is strictly concave. Thus, $\Delta_b(\gamma) > \xi_0(\gamma) - \theta_1 \rho(-\xi_0(\gamma))$. As for (b), notice now that $\xi_0(0) > 0$ since X_a second order stochastically dominates $2X_a$ and because ϕ is strictly concave. Hence, all terms on the right hand side of the previous inequality

are strictly positive and so $\Delta_b(0) > 0$, establishing (b). As for (c), note that by definition, $\xi_0(\bar{\gamma}) = 0$. Hence, $\Delta_b(\bar{\gamma}) > 0$.

As for (a), we have already established in the proof of Lemma 6 that $\lim_{\gamma \rightarrow \infty} \xi_0(\gamma) = -\infty$. Hence, it is sufficient to show that $\lim_{\gamma \rightarrow \infty} \xi_1(\gamma) = -\infty$. To see this notice that we have shown in part (a) of the proof of Lemma 6 that $E \left[\phi \left(X_a + \gamma \widehat{X}_a \right) \right]$ converges monotonically to $-\infty$ as $\gamma \rightarrow \infty$. Thus, since ρ is increasing, $\rho \left(\phi(2x_a) - E \left[\phi \left(X_a + \gamma \widehat{X}_a \right) \right] \right)$ converges monotonically to $-\infty$ for all x_a as $\gamma \rightarrow \infty$. Thus the integral over this expression converges monotonically to $-\infty$ by the Monotone Convergence Theorem. This establishes part (i).

Part (ii), which deals with the case, in which $s_a = b$, and the agent has chosen b in period 1, follows with identical arguments. \square

Proof of Lemma 8 Let $s_{y_1} \in Y_2$ be the agent's period 2 strategy after having chosen action y_1 in period 1. To show the claim, we have to show that for $\alpha \geq \widehat{\alpha}$, $s_a = y_2$ is a best response against $s_b = y_2$ and vice versa. We only show that $s_a = c$ is a best response against $s_b = c$. All the other cases follow with identical arguments.

So let $s_b = c$ and suppose the agent has chosen a in period 1. Then, the counterfactual self obtains $\phi(X_c)$ with probability α and $\phi(\widehat{X}_c)$ with probability $1 - \alpha$. Thus, the agent's reference point is $\phi(X_c)$ with probability α and $\phi(\widehat{X}_c)$ with probability $1 - \alpha$. Hence, his expected utility from choosing c is

$$\begin{aligned} u(c) &= E[\phi(X_c)] + \theta_1 E \left[\alpha \rho(\phi(X_c) - \phi(X_c)) + (1 - \alpha) \rho(\phi(X_c) - \phi(\widehat{X}_c)) \right] \\ &= E[\phi(X_c)] + \theta_1 (1 - \alpha) E \left[\rho(\phi(X_c) - \phi(\widehat{X}_c)) \right]. \end{aligned}$$

Likewise, his expected utility from choosing d is

$$\begin{aligned} u(d) &= E[\phi(X_d)] + \theta_1 E \left[\alpha \rho(\phi(X_d) - \phi(X_c)) + (1 - \alpha) \rho(\phi(X_d) - \phi(\widehat{X}_c)) \right] \\ &= E[\phi(X_d)] + \theta_1 E[\rho(\phi(X_d) - \phi(X_c))], \end{aligned}$$

where the second equality holds because X_c and \widehat{X}_c are identically distributed.

We want to show that $u(c)$ is larger than $u(d)$ for sufficiently large θ_1 and α . To see that is true, notice first that the $E[\phi(X_d)] \geq E[\phi(X_c)]$, because X_d second order stochastically dominates X_c by assumption. Thus, we need to show that the regret from c is smaller than that from d for sufficiently large α , that is

$$(1 - \alpha) E \left[\rho(\phi(X_c) - \phi(\widehat{X}_c)) \right] > E[\rho(\phi(X_d) - \phi(X_c))] \quad (6)$$

for sufficiently large α . (We can then choose θ_1 sufficiently large to override material incentives.) To (6), notice first that $\rho(E[\phi(X_d) - \phi(X_c)])$ is a strict upper bound for the right

hand side of (6) by Jensen's inequality and because ρ is strictly concave. Moreover, because X_d second order stochastically dominates X_c , this upper bound is negative, and hence the right hand side of (6) is strictly negative. Observe now that the left hand side (6) converges to 0 as α converges to 1. Thus, (6) holds for sufficiently large α , and this completes the proof. \square

Proof of Proposition 4 Suppose first the counterfactual self chooses $e = 0$. This gives rise to the reference point $R_1 = \phi(0)$. Thus, the agent's expected utility from choosing $e = 0$ and $e = 1$, respectively, is

$$\begin{aligned} u(0, 0) &= \phi(0) + \theta_1 \rho(0) = \phi(0), \\ u(1, 0) &= \phi\left(\frac{1}{2} - k\right) + \theta_1 \left(\frac{1}{2} \rho(\phi(1 - k) - \phi(0)) + \frac{1}{2} \rho(\phi(-k) - \phi(0)) \right). \end{aligned}$$

Hence,

$$\Delta_0 = u(1, 0) - u(0, 0) = \phi\left(\frac{1}{2} - k\right) - \phi(0) + \theta_1 \left(\frac{1}{2} \rho(\phi(1 - k) - \phi(0)) + \frac{1}{2} \rho(\phi(-k) - \phi(0)) \right).$$

Suppose next the counterfactual self chooses $e = 1$. If the agent chooses $e = 0$, then his reference point depends on the true state that is revealed to the agent. If he learns that $\omega = 1$, his reference point is $R_1 = \phi(1 - k)$, while if he learns that $\omega = 0$, his reference point is $R_1 = \phi(-k)$. Hence, his expected utility is

$$u(0, 1) = \phi(0) + \theta_1 \left(\frac{1}{2} \rho(\phi(0) - \phi(1 - k)) + \frac{1}{2} \rho(\phi(0) - \phi(-k)) \right).$$

If the agent chooses $e = 1$, then he matches exactly his counterfactual self's payoff and thus receives expected utility

$$u(1, 1) = \phi\left(\frac{1}{2} - k\right) + \theta_1 \rho(0) = \phi\left(\frac{1}{2} - k\right)$$

Hence,

$$\Delta_1 = u(1, 1) - u(0, 1) = \phi\left(\frac{1}{2} - k\right) - \phi(0) - \theta_1 \left(\frac{1}{2} \rho(\phi(0) - \phi(1 - k)) + \frac{1}{2} \rho(\phi(0) - \phi(-k)) \right).$$

We want to show that $\Delta_1 - \Delta_0 \geq 0$. From the above expressions, this difference can be written as

$$\begin{aligned} \Delta_1 - \Delta_0 &= -\frac{1}{2} \theta_1 \{ [\rho(\phi(0) - \phi(1 - k)) + \rho(\phi(1 - k) - \phi(0))] + \\ &\quad + [\rho(\phi(0) - \phi(-k)) + \rho(\phi(-k) - \phi(0))] \}. \end{aligned}$$

But because ρ is concave, the terms in the square brackets are non-positive, and this implies the claim. \square

References

- Aspinwall, L. G., L. Richter, and R. R. Hoffman (2001). "Understanding how optimism "works": An examination of optimists' adaptive moderation of belief and behavior." In *Optimism and Pessimism: Theory, Research, and Practice*, edited by E. C. Chang, Washington, American Psychological Association.
- Bell, D. E. (1982). "Regret in decision making under uncertainty." *Operations Research*, 30: 961-981.
- Bell, D. E. (1983). "Risk premiums for decision regret." *Management Science*, 29: 1156-1166.
- Bernheim, D. B. (1984). "Rationalizable strategic behavior." *Econometrica*, 52 (4): 1007-1028.
- Camerer, C. F. and R. A. Weber (1999). "The econometrics and behavioral economics of escalation of commitment: a reexamination of Staw and Hoang's NBA data." *Journal of Economic Behavior and Organisation*, 39: 59-82.
- Camille, N., G. Coricelli, J. Sallet, P. Pradat-Diehl, J.-R. Duhamel, and A. Sirigu. (2004): "The involvement of the orbitofrontal cortex in the experience of regret." *Science*, 304: 1167-1170.
- Caplin, A. and J. Leahy (2001). "Psychological expected utility theory and anticipatory feelings." *Quarterly Journal of Economics*, 116: 55-79.
- Eyster (2002). "A taste for consistency." *mimeo*.
- Festinger, L. and E. Walster (1964). "Post-decision regret and decision reversals." In Festinger, L. (ed). *Conflict, Decision Making, and Dissonance*. Stanford University Press, Stanford, California.
- Geanakoplos, J., D. Pearce, and E. Stacchetti (1989). "Psychological games and sequential rationality." *Games and Economic Behavior*, 1: 60-79.
- Harless, D. (1992). "Actions versus prospects: the effect of problem representation on regret." *American Economic Review*, 82: 634-649.
- Heidhues, P. and B. Koszegi (2004). "The impact of consumer loss aversion on pricing." *mimeo*.

- Inman, J. J., J. S. Dyer, and J. Jia (1997). "A generalized utility model of disappointment and regret effects on post-choice valuation." *Marketing Science*, 16: 97-111.
- Irons, B. and C. Hepburn (2003). "Regret Theory and the tyranny of choice." *mimeo*.
- Josephs, R. A., R. P. Larrick, C. M. Steele, and R. E. Nisbett (1993): "Protecting the self from the negative consequences of risky decisions." *Journal of Personality and Social Psychology*, 1992, 62 (1): 26-37.
- Kahneman, D. and A. Tversky (1979). "Prospect theory: An analysis of decision under risk." *Econometrica*, 47: 263-291.
- Kahneman, A. and D. T. Miller (1986). "Norm theory: comparing reality to its alternatives." *Psychological Review*, 93 (2): 136-153.
- Kardes, F. R. (1994). "Consumer judgement and decision processes." In Wyer, R. S. and T. C. Srull (eds). *Handbook of Social Cognition*, Vol 2 (2), Hillsdale, NJ, Lawrence Erlbaum Associates.
- Kelsey, D. and A. Schepanski (1991). "Regret and disappointment in taxpayer reporting decisions: an experimental study," *Journal of Behavioral Decision Making*: 4, 33-53.
- Koszegi, B. (2000a). "Ego utility and information acquisition." *mimeo*.
- Koszegi, B. (2000b). "Ego utility, overconfidence and task choice." *mimeo*.
- Koszegi, B. (2004). "Utility from anticipation and personal equilibrium." *mimeo*.
- Koszegi, B, and M. Rabin (2004). "A model of reference dependent preferences." *mimeo*.
- Larrick, R. P. and T. J. Boles (1995). "Avoiding regret in decisions with feedback: A negotiation example." *Organizational Behavior and Human Decision Processes*, 63(1): 87-97.
- Loomes, G. (1988a). "Further evidence of the impact of regret and disappointment in choice under uncertainty." *Economica*, 55: 47-62.
- Loomes G. (1988b). "When actions speak louder than prospects." *American Economic Review*, 78: 463-470.
- Loomes, G. and R. Sugden (1982). "Regret theory: An alternative approach to rational choice under uncertainty." *The Economic Journal*, 92: 805-824.

- Loomes, G. and R. Sugden (1987a). "Some implications of a more general form of regret theory." *Journal of Economic Theory*, 41: 270-87.
- Loomes, G. and R. Sugden (1987b). "Testing for regret and disappointment in choice under uncertainty." *The Economic Journal*, 97 (Supplement): 118-129.
- Loomes, G., C. Starmer and R. Sugden (1991). "Observing violations of transitivity by experimental methods." *Econometrica*, 59(2): 425-439.
- Loomes, G., C. Starmer and R. Sugden (1992). "Are preferences monotonic? Testing some predictions of regret theory." *Economica*, 59: 17-33.
- Markus, H. and P. Nurius (1986). "Possible Selves." *American Psychologist*, 41(9): 954-969.
- Mellers, B. (2000). "Choice and the relative pleasure of consequences." *Psychological Bulletin*, 126: 910-924.
- Mellers, B., A. Schwartz and I. Ritov (1999). "Emotion-based choice." *Journal of Experimental Psychology: General*, 128: 332-345.
- Mellers, B. and A. P. McGraw (2001). "Anticipated emotions as guides to choice." *Current Directions in Psychological Science*, 10(6): 210-214.
- Prendergast, C. and L. Stole (1996). "Impetuous youngsters and jaded old timers: acquiring a reputation for learning." *Journal of Political Economy*, 104(6): 1105-1134.
- Rabin, M. (1998). "Psychology and Economics." *Journal of Economic Literature*, 36: 11-46.
- Ritov, I. (1996). "Probability of regret: anticipation of uncertainty resolution in choice." *Organizational Behavior and Human Decision Processes*, 66(2): 228-236.
- Roese, N. J. (1997). "Counterfactual thinking." *Psychological Bulletin*, 121(1): 133-148.
- Simonson, I. (1992). "The influence of anticipating regret and responsibility on purchase decision." *Journal of Consumer Research*, 19:105-118.
- Starmer, C., and R. Sugden (1993). "Testing for juxtaposition and event-splitting effects." *Journal of Risk and Uncertainty*, 6: 235-254.
- Starmer, C., and R. Sugden (1998). "Testing alternative explanations for cyclical choice." *Economica*, 65: 347-361.

- Taylor, S. E. and J. D. Brown (1988). "Illusion and well-being: a social psychological perspective on mental health." *Psychological Bulletin*, 103(2): 193-210.
- Taylor, K. A. (1997). "A regret theory approach to assessing consumer satisfaction." *Marketing Letters*, 8(2): 229-238.
- Tsiros, M. and V. Mittal: (2000). "Regret: A model of its antecedents and consequences in consumer decision making." *Journal of Consumer Research*, 26: 401 - 417.
- Tykocinski, O. E. and T. S. Pittman (1998). "The consequences of doing nothing: Inaction inertia as avoidance of anticipated counterfactual regret." *Journal of Personality and Social Psychology*, 75(3): 607-616.
- Yariv, L. (2005). "I'll see it when I believe it: A simple model of cognitive consistency." *mimeo*.
- Zeelenberg, M. (1999). "Anticipated regret, expected feedback, and behavioral decision making." *Journal of Behavioral Decision Making*, 12: 93-106.
- Zeelenberg, M. (2004). "Consequences or regret aversion in real life: the case of the Dutch postcode lottery." *Organizational Behavior and Human Decision Processes*, 93: 155-168.
- Zeelenberg, M., J. Beattie, J. Van Der Pligt, and N. K. De Vries (1996). "Consequences of regret aversion: Effects of expected feedback on risky decision making." *Organizational Behavior and Human Decision Processes*, 65(2): 148-158.
- Zeelenberg, M., and J. Beattie (1997). "Consequences of regret aversion 2: Additional evidence of the effects of feedback on decision making." *Organizational Behavior and Human Decision Processes*, 72(1): 63-78.