



Working Papers

www.cesifo.org/wp

A Model of Reference-Dependent Belief Updating

Johannes K. Maier
Clemens König

CESIFO WORKING PAPER NO. 6156
CATEGORY 13: BEHAVIOURAL ECONOMICS
OCTOBER 2016

An electronic version of the paper may be downloaded

- *from the SSRN website:* www.SSRN.com
- *from the RePEc website:* www.RePEc.org
- *from the CESifo website:* www.CESifo-group.org/wp

ISSN 2364-1428

A Model of Reference-Dependent Belief Updating

Abstract

We propose a model of instrumental belief choice under loss aversion. When new information arrives, an agent is prompted to abandon her prior. However, potential posteriors may induce her to take actions that generate a lower utility in some states than actions induced by her prior. These losses loom larger than gains in other states in which potential posteriors lead to larger utility than the prior. In choosing her belief, the agent optimally trades off these psychological gains and losses against the belief's objective performance. Consistent with empirical evidence, belief updating in this model is conservative and can be asymmetric, so that bad news is updated more conservatively than good news. These updating biases generate prior-dependent information preferences, such that agents may avoid information when being unconfident but are information seeking otherwise. Because belief updating depends on the decision problem in which new information is going to be used, the model allows us to explore how these predictions change with either individual preferences or the choice context.

JEL-Codes: D030, D810, D830, D840, G020.

Keywords: belief choice, non-Bayesian updating, reference dependence, loss aversion, regret, conservatism, signal valence, information preferences, overconfidence.

Johannes K. Maier
Department of Economics
Seminar for Economic Theory
University of Munich
Ludwigstr. 28 Rgb.
Germany – 80539 Munich
johannes.maier@econ.lmu.de

Clemens König
NERA Economic Consulting
Marble Arch House
66 Seymour Street
United Kingdom – London W1H 5BT
clemens.koenig@gmail.com

This Version: October 2016

The authors are grateful to Tore Ellingsen, Moshe Hoffman, Martin G. Kocher, Botond Köszegi, Augustin Landier, Peter Schwardmann, Andrei Shleifer, Adam Szeidl, Bertrand Villeneuve, Florian Zimmermann, and especially to Lukas Buchheim, Dominik Fischer, Fabian Herweg, Takeshi Murooka, Klaus M. Schmidt, and Marco Schwarz for helpful comments. We also thank seminar participants at Humboldt University Berlin, University of Bayreuth, University of Munich, and at the following meetings for valuable comments and questions: CESifo Area Conference on Behavioral Economics, CEAR/MRIC Behavioral Insurance Workshop, Thurgau Experimental Economics Meeting, Russell Sage Foundation ECBE Conference, SFB/TR15Workshop, German Economists Abroad Meeting, CESifo Area Conference on Applied Microeconomics, and EEA/ESEM Meeting.

“First, persons not only have preferences over states of the world, but also over their beliefs about the state of the world. Second, persons have some control over their beliefs; not only are people able to exercise some choice about belief given available information, they can also manipulate their own beliefs by selecting sources of information likely to confirm ‘desired’ beliefs.”

— Akerlof and Dickens, 1982, p.307

1 Introduction

A fundamental issue in economics is how people incorporate new information about the uncertain world around them into their decision making. The standard assumption in this regard is that people are Bayesians, who hold a prior belief about the world and use Bayes’ Rule to update it in response to new information. Yet, this mechanical model of belief formation has been increasingly called into question. Empirical evidence suggests that beliefs are often revised *conservatively*, i.e., change in the same direction as under Bayesian updating, but less so. Also, belief updating is often *asymmetric* with “bad news” typically giving rise to more conservatism than “good news.” When information acquisition is endogenous, there is growing evidence that people often have a negative willingness to pay for additional information, i.e., they are information averse, if their confidence in the “good state” is low. These inconsistencies with the Bayesian model have important consequences, e.g., for financial markets, employment relationships, and educational choices.

This paper proposes a model of *reference-dependent* belief formation that provides the first unified explanation for these phenomena. It explains both conservative and asymmetric updating and – through the anticipation of such biased updating – allows for a novel form of *prior-dependent* information aversion in certain well-defined circumstances that are consistent with the empirical evidence. Our model builds on Akerlof and Dickens (1982) by assuming that beliefs are *objects of choice* and that this choice is affected by the preferences of the decision maker. However, contrary to existing models of belief choice, our choice procedure remains entirely *instrumental*, meaning that the value of a given belief depends solely on the action(s) that the belief implements.

Two components determine the instrumental value of beliefs: *objective performance* and *loss aversion*. Objective performance refers to the expected utility of a belief’s implemented action given the Bayesian posterior containing the new information. The impact of loss aversion is more subtle: New information prompts the agent to abandon her status quo belief. Yet, giving up one’s prior cannot be preferable in all states of the world. Any new belief inducing an action different from the prior-induced *reference action* has to perform worse in at least one state of the world. Otherwise, the reference action would be dominated and not rationalized by the prior in the first place. Thus, whenever new information invites the agent to adjust her belief *and* this new belief implements a different action than the prior, she faces the dilemma of there being at least one state where holding on to her prior would have been preferable. In such states, psychological losses are experienced, and these losses loom larger than the gains derived in states where the new belief is preferable. Our model builds on the idea that the avoidance of such losses is a second goal in belief choice next to objective performance, with the ultimately chosen belief representing the best compromise between these two objectives.¹

¹An alternative interpretation for the second goal in belief choice is *prior regret* avoidance. In states where holding on to one’s prior would have been preferable, there is scope for (ex-post) regret from having abandoned one’s prior. We return to this issue in our discussion of the relevant literature.

To illustrate our approach, consider the well-known “Monty Hall” problem: A decision maker D participating in a game show has a choice of opening one of three doors behind one of which there is a prize. After D has made a preliminary choice, the show’s host opens one of the remaining doors, but in doing so avoids the door concealing the prize. Then, D is given the opportunity either to stick to her preliminary choice or to switch doors. It is well-known that switching maximizes the chance of winning. Yet, experimental subjects show a strong propensity to stick to their original choice, even after receiving ample feedback for learning the rational decision (e.g., Friedman, 1998; Palacios-Huerta, 2003). In terms of our model, we can explain this propensity by the trade-off outlined above. From the viewpoint of objective performance, those beliefs about the prize’s location are best that cause D to switch doors since switching yields the highest expected utility given the correct Bayesian posterior. Loss aversion points into a different direction: All beliefs causing D to switch doors entail a loss in the state of the world where the prize is behind the preliminarily chosen door (her reference action). In contrast, all beliefs maintaining D’s preliminary choice cause no such loss feelings and are thus superior from the point of view of loss avoidance. All in all, if D pays sufficient heed to her expected losses from switching (where the expectation is calculated using the Bayesian posterior) all beliefs causing her to stick to her preliminary choice are optimal despite their lower objective performance.²

The example raises the question why we assume that belief choice rather than action choice is affected by (status quo) loss aversion.³ Three main reasons favor our approach: First, we are interested in explaining how and when *beliefs* deviate from the Bayesian benchmark, for which the action choice model is not suited. Second, the action choice model faces the conceptual problem that it may lead to choices that are first-order stochastically dominated given the decision maker’s belief.⁴ The belief choice model avoids this problem because the predicted action is always optimal given the chosen belief. Third and most importantly, because agents will acquire different information, the two models generally imply different *behavior*.⁵ In particular, our model can explain why people sometimes ignore costless information, which cannot be explained by a model of action choice.

After discussing the related theoretical literature in Section 2, we formally present our model in Section 3. In this section, we also investigate three “boundary” properties of our model. First, without loss aversion our model reduces to standard Bayesian decision making. In the other extreme, under “strong” loss aversion, our model generates a *preference for consistency*, i.e., the tendency to replicate one’s previous action choice regardless of new information (e.g., Cialdini,

²Of course, such non-switching behavior may also be due to *unmotivated* misconceptions on the posterior probability, e.g., from applying Laplace’s principle of insufficient reason to both the prior and the posterior. Still, the example nicely illustrates the trade-off our model captures. Moreover, there is evidence in support of our explanation: Granberg and Brown (1995) and Gilovich et al. (1995) provide evidence that counterfactual thoughts in the Monty Hall problem predominantly arise for “switch losses” rather than “stick losses”. And Petrocelli and Harris (2011) show that this “omission bias” inhibits subjects’ learning in the Monty Hall problem.

³The same explanation for people’s reluctance to switch doors could be generated by a model of action choice where preferences incorporate loss aversion relative to one’s previous action, but belief formation is entirely Bayesian.

⁴For example, as long as D’s prior confidence that the prize is behind her initially chosen door is somewhere below 50%, her Bayesian posterior after one of the other two doors has been opened will put more probability mass on the prize being behind the door that she did not choose. In these conditions, “switching” first-order stochastically dominates “sticking” given D’s posterior.

⁵Ambuehl and Li (2015) provide evidence that deviations of individuals’ valuation of information from the standard model are driven by the anticipation of non-Bayesian belief updating rather than non-standard preferences given the Bayesian posterior. This is consistent with our model where biased information preferences are *caused* by (the anticipation of) biased belief updating.

1984; Falk and Zimmermann, 2015). Finally, if all uncertainty is resolved before a belief is chosen, e.g., via a signal that is fully precise, our decision maker behaves like a Bayesian *irrespective* of her loss aversion. In terms of the Monty Hall problem, if the show’s host accidentally opens the door concealing the prize, it is optimal for D to “believe her eyes” and to choose the door where she perceives the prize to be. This intuitive result, which other models of belief choice need to assume, is a consequence of our instrumental approach.

The remainder of the paper focuses on “interior” properties of our model. In Section 4, we address two anomalies pertaining to *information processing* that have been widely documented by the empirical literature. We first show that our model generates conservatism, a long studied phenomenon in the psychology literature (building on Edwards, 1968). An agent in our model updates conservatively upon receiving contradicting news because she fears the sensation of loss from realizing that holding on to her prior belief would have led her to take preferable actions. The size of the loss is determined by her state- as well as action-contingent utility levels, and, as a result, her conservatism is affected by these utilities. This illustrates that updating biases in our model are not due to cognitive errors, but rather motivated by the decision problem in which information is going to be used and thus context-dependent. While being consistent with recent experiments revisiting conservatism (Möbius et al., 2014; Ambuehl and Li, 2015), this differentiates our approach from models of mechanically erroneous under-updating (e.g., Barberis et al., 1998; Rabin and Schrag, 1999; Benjamin et al., 2016). In addition to context dependency, the instrumental nature of our approach has another distinct implication: individuals who hold identical prior beliefs will update differently if their prior-induced actions differ (e.g., due to differing risk attitudes). Recent studies show that, when controlling for prior beliefs, updating is conservative for news that contradict prior actions (Kuhnen and Knutson, 2011; Kuhnen et al., 2015), providing evidence for the mechanism that is at the heart of our model.

Another phenomenon at odds with mechanical updating models (Bayesian or otherwise) is the fact that signal valence matters for information processing (Eil and Rao, 2011; Möbius et al., 2014; Wiswall and Zafar, 2015; Sharot et al., 2011). A decision maker in our model updates good news less conservatively than bad news whenever the utility loss from the wrong action is larger in the good than bad state. For example, an investor becomes more prone to such asymmetric updating with the “upside” of her investment (relative to not investing) exceeding its “downside.” By contrast, in a neutral choice context – where states are equally good and guessing correctly is the only concern (as in the Monty Hall problem) – not matching the state leads to the same utility losses in all states. In this case, our model predicts symmetric, i.e., equally conservative, updating. Such differences in updating behavior not only accord with the evidence, but also provide new testable predictions in which environments we should expect more, no, or a reversal of asymmetric updating. Also, our model sheds new light on how asymmetry interacts with risk attitudes: because risk aversion decreases the difference between the investment’s upside and downside, the investor becomes more prone to asymmetric updating the more risk tolerant she is. As a result, our model suggests overconfidence being negatively related to risk aversion.

Moving beyond the processing of exogenously given information, in Section 5 we make information acquisition endogenous and analyze agents’ *information preferences*. In contexts where news is intrinsically good or bad, recent evidence suggests that individuals are information seeking when being confident, but avoid information when being less confident in the good state (Eil

and Rao, 2011; Burks et al., 2013; Möbius et al., 2014; Karlsson et al., 2009). While such prior-dependent information preferences seem highly intuitive, they are at odds not only with Bayesian updating but, as shown by Eliaz and Spiegel (2006) or Spiegel (2008), also with models in which beliefs have direct utility value (e.g., Caplin and Leahy, 2001; Brunnermeier and Parker, 2005; Kőszegi, 2006a). Our instrumental model generates this type of information aversion whenever taking the inappropriate action in the good state is still better than taking the appropriate action in the bad state. These are settings in which people care strongly about the state of Nature, e.g., being attractive or intelligent, and no action can make the reverse state preferable.⁶ By contrast, in neutral choice contexts our decision maker is information seeking irrespective of her prior confidence level. Still, she differs from a Bayesian agent in a way consistent with recent evidence, e.g., underpays for valuable but overpays for less valuable information (Hoffman, 2014), or has a willingness to pay that responds too little to changes in the signal’s informativeness (Hoffman, 2014; Ambuehl and Li, 2015).

Section 6 is devoted to a robustness analysis: We let the action space be continuous rather than finite, which allows us to derive a unique point prediction for the agent’s optimal belief.⁷ All our results are shown to extend to this richer setting. We conclude in Section 7. Appendix A is concerned with the empirical implementation of our model: We first demonstrate that our approach is compatible with the (binomial) logistic regression model, meaning that optimal beliefs can be estimated using the standard econometric approach to a framework like ours. Secondly, we propose an experimental choice task to elicit the free parameter of our model (measuring loss aversion). Proofs are all gathered in Appendix B.

2 Related Theoretical Literature

Starting with Akerlof and Dickens (1982), a few papers build on the idea that beliefs are *objects of choice* (rather than being formed in the mechanical fashion envisaged by Bayes’ Rule). Next to objective performance, the target function for belief choice takes different forms: In Akerlof and Dickens (1982), workers want to minimize the “fear” of having an accident. Yariv (2005) defines a property of beliefs called “directional confidence,” meaning that the decision maker values increasing her level of confidence in the state of the world she regards as most likely. Brunnermeier and Parker (2005) take it to include “anticipatory utility,” which is defined as the expected utility of the action(s) implemented by each belief calculated using the belief itself (see also Landier, 2000). Applications of their “optimal expectations” framework are provided by Brunnermeier et al. (2007) or Oster et al. (2013), while Spiegel (2008) demonstrates the model’s inability to generate information aversion, a prediction different from our approach (we return to this issue in Section 5). Furthermore, all these models have a tendency to create extreme beliefs (or beliefs that are severely tilted towards one state) which makes them incompatible with conservative updating.⁸ Also, beliefs enter the target function directly in these models. In contrast, our model of belief choice is the first where the choice procedure is fully instrumental in that beliefs enter the

⁶Even if one has chosen not to obtain higher education, higher ability is still likely to be preferred. And even if one has not dared to approach the (wo)man of one’s dreams, a preference for higher attractiveness seems plausible.

⁷Due to the instrumental nature of our model, any two beliefs inducing the same action are equally valued. Hence, when actions are finite, there may be more than one optimal belief.

⁸The model of “wishful thinking” (Mayraz, 2014) also creates such tilted beliefs, however, mechanically: an exogenous parameter determines by how much agents distort the Bayesian posterior to be either pessimistic or optimistic.

target function for belief selection only via the action(s) that they implement. We take this to be a parsimonious specification where departures from the standard approach have been minimized.⁹

A second strand of literature does not view beliefs as objects of choice, but develops alternative belief *mechanics* intended to replace Bayes' Rule. Some of these approaches generate conservatism, but they all face difficulties in accounting for asymmetry: Rabin and Schrag (1999) propose a model of "confirmatory bias," where signals may be erroneously perceived as confirming the prior whenever prior and signal point into opposite directions.¹⁰ Benjamin et al. (2016) model a "non-belief in the law of large numbers," which leads to under-inference as people believe that proportions of binary signals depart from the population mean even in very large samples.¹¹ In the model of Barberis et al. (1998), conservatism is caused by investors' irrational perception that earnings are mean-reverting.¹² In the models of Epstein et al. (2010) or Bouchaud et al. (2016), the decision maker processes new information in such a way that her adopted belief is a weighted average of the Bayesian posterior and her prior.¹³ Our model can be viewed as providing a micro-foundation for the decision maker placing positive weight on her prior. In general, however, the foundational idea of our model is that information processing cannot be dissociated from the decision problem in which the information is (likely to be) employed. This differentiates our approach from mechanical, "context-free" models of belief revision where the standard separation between beliefs and utility is preserved.¹⁴

Thirdly, several papers examine the idea that beliefs – formed according to Bayes' Rule – enter the utility function for action choice directly. Kőszegi (2006a) studies "ego utility," where one's belief about own ability is a direct carrier of utility (see also Weinberg, 2009). In the model of "psychological expected utility" introduced by Caplin and Leahy (2001), agents derive (an extra) anticipatory utility equal to their currently expected future utility. The model has found some interesting (mostly health-related) applications (Kőszegi, 2003, 2006b; Caplin and Leahy, 2004; Barigozzi and Levaggi, 2010). In both approaches beliefs can only be distorted through acquired information. Applications have therefore focused on anomalous attitudes to information, in particular, conditions under which less accurate information is preferred.¹⁵ Yet, results differ starkly com-

⁹The belief choice models of Gollier and Muermann (2010) and Macera (2014) depart from an instrumental nature of beliefs even further as *all* elements in the trade-off for belief selection are essentially non-instrumental.

¹⁰For other forms of confirmatory bias see Schwartzstein (2014) or Fryer et al. (2015).

¹¹For samples of size one, the model reduces to Bayesian updating. When extensively reviewing the evidence on conservatism the authors note, however, that conservative updating is also common in samples of size one, suggesting that there may be an additional source for conservatism.

¹²As a result, prices underreact to news (leading to "drift" and "momentum") because investors wrongly believe that changes are temporary. However, after observing a series of similar news investors also wrongly believe that earnings are trending and extrapolate this trend. Hence, prices overreact in the long term leading to reversal of returns (e.g., DeBondt and Thaler, 1985). Unlike Barberis et al. (1998), our model is static and not tailor-made to explain the simultaneous (shorter-term) underreaction and (long-term) overreaction of stock markets. We are therefore unable to say anything about these interesting long-term dynamics. Gennaioli et al. (2015) sketch a model based on representativeness (Kahneman and Tversky, 1972) and Mullainathan (2002) uses a memory-based approach to explain simultaneous over- and underreaction to news.

¹³Epstein et al. (2010) investigate what under- and overreaction to news implies for belief accuracy in the long run. Drawing from the macro literature on sticky expectations, Bouchaud et al. (2016) show that underreaction not only explains momentum but also the "quality" anomaly.

¹⁴Other mechanical models lead to over- rather than under-inference and capture biases such as a "belief in the law of small numbers" (Tversky and Kahneman, 1971; Rabin, 2002; Rabin and Vayanos, 2010), "base rate neglect" (Kahneman and Tversky, 1973; Grether, 1980; Bodoh-Creed, 2009), and further biases caused by representativeness (Kahneman and Tversky, 1972; Gennaioli and Shleifer, 2010; Bordalo et al., Forthcoming, 2016).

¹⁵Another model that has been applied to such contexts is Kőszegi and Rabin (2009), where utility depends on *changes* in (rational) beliefs rather than beliefs themselves because belief changes inform about consumption.

pared to ours. According to Eliaz and Spiegel (2006) or Spiegel (2008), psychological expected utility fails to account for realistic *prior-dependent* attitudes to information, where agents seek information when being confident but avoid it when being not so confident.¹⁶ In the model of ego utility, information aversion is depending on the prior but in the opposite direction: unconfident agents seek information and confident ones avoid it in order to prevent a potential downward revision of their belief (see Burks et al., 2013). Moreover, since these models assume standard belief formation, they have a harder time to account for the evidence on non-standard belief updating.

Fourth, several papers build on the idea that an agent's multiple selves play intra-personal games. On the demand side, i.e., *why* people may hold distorted beliefs, it has been argued that agents have limited willpower (Carrillo and Mariotti, 2000; Bénabou and Tirole, 2002, 2004, 2006, 2011) or anticipatory utility (Bénabou and Tirole, 2011; Bénabou, 2013). On the supply side, i.e., *how* agents manage to hold distorted beliefs, agents can strategically ignore information (Carrillo and Mariotti, 2000), have selective memory (Bénabou and Tirole, 2002, 2006; Bénabou, 2013), or generate their own signals, i.e., "self-signal" (Bénabou and Tirole, 2004, 2011). Our approach differs from these models in several respects: Instead of determining equilibrium play among an agent's multiple selves, our results are rather the outcome of a non-strategic individual choice problem. Concerning the supply side, we do not restrict which beliefs may be chosen in any way. Also, our demand side is neither specifically targeted to decision problems in which limited willpower causes frictions nor does it confine the instrumentality of beliefs by giving them direct utility value.

Finally, building on the idea of prior-induced loss aversion, our model situates itself within the large literature on reference-dependent choice. Several formalizations of reference dependence exist.¹⁷ However, providing the concept's first application to belief formation, we use a reference point not previously considered in the literature: the action(s) induced by one's prior. We view the prior as *the* natural and salient reference point for belief updating because the prior can be thought of as the *status quo* belief to which individuals have developed some kind of attachment and new information invites them abandoning it. This relates our model to the literature on the status quo bias (Samuelson and Zeckhauser, 1988). Loss aversion (see the literature building on Kahneman and Tversky, 1979; Tversky and Kahneman, 1991) has been used to explain this bias since, with the reference point being status quo consumption, one dislikes moving away from it as induced consumption losses loom larger than same-sized consumption gains.¹⁸ Another possible explanation for the status quo bias roots in the literature on regret. Although in models of regret (Loomes and Sugden, 1982; Bell, 1982) the reference point is not fixed but rather the alternative not taken by the decision maker, a common finding in the empirical literature on regret and counter-

¹⁶More specifically, Eliaz and Spiegel (2006) show that if an agent has an "extreme" prior and ranks a noiseless signal either above or below all other signals, she will do so for all priors. This is different in our model: When the prior is close to 1, the noiseless signal is preferred over all noisy signals. Yet, there exist lower priors where this is not necessarily the case.

¹⁷All existing models of reference dependence cover action choice. In contrast, for a given belief, action choice in our model is reference-*independent*. It is only belief choice that is reference-dependent. Although it may well be plausible that action choice is reference-dependent as well (e.g., as in Kőszegi and Rabin, 2006, 2007), we do not integrate it into our model in order to identify the pure effects of reference-dependent belief choice.

¹⁸Loss aversion further relates our model to the literature on the endowment effect (Thaler, 1980; Kahneman et al., 1990). It implies, however, that people are reluctant to part with their *beliefs* rather than *possessions*, providing a link to an idea of Abelson (1986): "I have already tipped it off in the title, 'Beliefs are like possessions.' Virtually everyone has had a great deal of experience dealing with possessions, and many orientations toward possessions can be generalized to belief objects" (p.229).

factual thinking – sometimes called the “omission bias” – is that people experience considerably more regret from action than from inaction (e.g., Kahneman and Tversky, 1982; Kahneman and Miller, 1986; Landman, 1987; Gleicher et al., 1990; Ritov and Baron, 1990; Gilovich and Medvec, 1995), a finding also supported by neurological evidence (e.g., Nicolle et al., 2011). In the extreme case where regret is felt only towards action, the anticipation of ex-post regret results in status quo loss aversion (regret also looms larger than rejoice).¹⁹ Similarly, regret can provide an intuition for the mechanics present in our model. People anticipate that they will (ex post) regret having abandoned their prior in response to new information if it turns out that their prior had been superior. Thus, they feel regret from action, but not from the inaction of having not updated accordingly. Our model can therefore also be interpreted as capturing “prior regret.”

3 The Model

The setup is the following. The decision-maker (D) first receives an informative signal s from some finite space S . The signal allows D to make inferences about the prevailing state of Nature $x \in X$. The space X is finite as well. After receiving her signal, D must take an action $a \in A$. The space A may be finite or continuous, and we will consider both specifications below. In general, the utility action $a \in A$ generates if the state of Nature is $x \in X$ is given by $u(x, a)$. Yet, instead of choosing directly from A , we model D as choosing from the set of possible beliefs about the state of Nature. These beliefs are collected in Q and contain both the prior p and the Bayesian posterior $\tilde{q}(p, s)$, i.e., the probability distribution on X implied by Bayes’ Rule as applied to p and s . The belief that D chooses then determines her action.²⁰

We make a few basic assumptions, whose motivation is technical. Firstly, the set Q of admissible beliefs about the state of Nature only contains probability distributions with full support on X .²¹ Secondly, and if not stated otherwise, signals are not fully revealing, i.e., we have $\Pr(s|x) \in (0, 1)$ for every $x \in X$ and $s \in S$. These two assumptions imply that also all updated beliefs have full support on X . And finally, there is “scope for loss aversion”: there exists at least one prior belief $\bar{p} \in Q$ and signal $\bar{s} \in S$ such that we do *not* have $A^*(\bar{p}) \subseteq A^*(\tilde{q}(\bar{p}, \bar{s}))$, where $A^*(f)$ are the optimal actions given some belief $f \in Q$, i.e., $A^*(f) = \arg \max_{a \in A} \sum_x f(x) u(x, a)$. Intuitively, scope for loss aversion means that there exists at least one prior belief and signal such that there is some tension at the level of implemented actions between adopting the Bayesian posterior and sticking to one’s prior. Imposing scope for loss aversion thus only rules out uninteresting cases where the trade-off we aim to capture in this model is absent.

3.1 Belief Choice

We next introduce the belief choice problem. Suppose that D holds the prior belief $p \in Q$ and has received the signal $s \in S$. In the following, we present how she then chooses her belief $q \in Q$.

¹⁹Interestingly, the importance of the status quo has already been recognized by Bell (1982): “The level of regret felt may sometimes be related to the original status quo no matter what the outcome of foregone alternatives” (p.980).

²⁰While we assume that information arrives and is exploited at most once, there are scenarios where a repeated setup seems more plausible. In such a scenario, the present framework remains valid only if agents are myopic or future information comes as a surprise. The case where agents foresee the future is addressed in another paper, in which we will explore the dynamics of reference-dependent belief updating.

²¹While this assumption is not needed for the chosen (posterior) belief, it is required for the prior.

First, notice that D 's choice of belief q pins down her subsequent action choice, the link between the two being provided by standard expected utility maximization. Formally, let

$$A^*(q) = \arg \max_{a \in A} \sum_x q(x) u(x, a)$$

denote the set of D 's preferred actions given q . We further denote by $G(A^*, q)$ the cumulative distribution function resulting from uniform randomization among all preferred actions, i.e., among all $a \in A^*(q)$. Usually, $A^*(q)$ will be a singleton, so that the single most preferred action is chosen with certainty, and $G(A^*, q)$ is degenerate. It is only when q makes D indifferent between various (or a range of) actions that she uniformly randomizes among them. We can now define the (expected) utility induced by q in state x as

$$U(x, q) = \int_{a \in A^*(q)} u(x, a) dG(A^*, q).$$

As mentioned before, the payoff of D 's belief choice q has two components: *objective performance* – which refers to the expected utility of q given the correct Bayesian posterior – and *loss aversion* – referring to the extent to which q minimizes D 's expected (net) loss relative to her prior. Capturing both components, the target function for D 's belief choice $q \in Q$ is given by

$$V(q|p, s) = \sum_x \Pr(x|p, s) \left[U(x, q) + r(U(x, q) - U(x, p)) \right],$$

where $\Pr(x|p, s)$ is the probability assigned to state x by the Bayesian posterior $\tilde{q}(p, s)$. The first part of the target function, namely $\sum_x \Pr(x|p, s) U(x, q)$, represents the expected utility of q under the Bayesian posterior and therefore the objective performance of q . The second part, $\sum_x \Pr(x|p, s) r(U(x, q) - U(x, p))$, expresses the expected gains and losses associated with q , calculated using again the Bayesian posterior and a reference point that is provided by D 's prior p .

With $y \in \mathbb{R}$ being the utility difference in state x , we assume that the gain-loss (or regret) function $r(y)$ is continuous for all y , sufficiently often differentiable for $y \neq 0$, and satisfies the following properties:

$$(i) r(y) = 0 \forall y \geq 0, \quad (ii) \lim_{y \rightarrow 0} r'(-|y|) \equiv \lambda \geq 0, \quad \text{and} \quad (iii) r''(y) \leq (=) 0 \forall y < 0 \text{ if } \lambda > (=) 0.$$

Condition (i) assumes that gains are truncated, so that $r(y)$ in the negative domain can be interpreted as “net loss utility” function capturing the sum over utilities of a loss and a same-sized gain.²² According to Condition (ii), $\lambda > 0$ induces a “kink” at the reference point, often referred to as “loss aversion in the small.” Moreover, if $\lambda > 0$, Condition (iii) implies that $r'(y) > 0$ for all $y < 0$, often referred to as “loss aversion in the large” and meaning that marginal utility of a loss exceeds marginal utility of a comparable gain (Bowman et al., 1999). This difference either

²²Precisely, for $y < 0$, $r(y) \equiv \mu(y) + \mu(-y)$ where $\mu(\cdot)$ is a conventional gain-loss function with the usual properties (see, e.g., Bowman et al., 1999). We impose Condition (i) for simplicity. Including gain feelings would not alter the conclusions from our model as long as losses loom larger than gains. In the literature on reference-dependent action choice, several papers ignore gains focusing on losses only (e.g., Herweg and Mierendorff, 2013; Herweg and Schmidt, 2015). We follow this literature, but allow for non-linear net loss utility.

increases ($r''(y) < 0 \forall y < 0$) or stays constant ($r''(y) = 0 \forall y < 0$) for larger losses.²³ $\lambda = 0$ refers to the case of no loss aversion, because Conditions (i) to (iii) imply that $r(y) = 0$ for all y in this case. On the other hand, if loss aversion is present and $\lambda > 0$, Conditions (i) to (iii) ensure that $r(y) < (=) 0$ for all $y < (\geq) 0$. Thus, D experiences a loss from q in a given state x if and only if $U(x, q)$ falls short of $U(x, p)$, which is the utility generated in x by her prior p and is analogously defined as $U(x, q)$ above. Clearly, setting $q = p$ yields no loss in every state, while other beliefs may give rise to losses.

How much D dislikes a specific loss $y < 0$, i.e., how negative $r(y)$ is, is generally determined by both the kink of $r(\cdot)$ at the reference point and the curvature of $r(\cdot)$ in the loss domain. Sometimes, however, it is useful to consider a special case of Condition (iii) in which not only the presence but also the *level* of loss aversion is determined by λ only:

Assumption 1 (Piecewise Linear Gain-Loss Utility) $r''(y) = 0 \forall y < 0$ if $\lambda \geq 0$.

Whenever Assumption 1 holds, without loss of generality we let $r(y) = \lambda y$ if $y < 0$ and $r(y) = 0$ otherwise, so that λ measures the strengths of loss aversion.

The optimal belief represents the best compromise between objective performance and loss aversion. Formally, the set of optimal beliefs, given prior p and signal s , is given by

$$Q^*(p, s) = \arg \max_{q \in Q} V(q | p, s).$$

Because D's belief choice is purely instrumental, $Q^*(p, s)$ need not be a singleton. Under a finite action space (see Sections 4 and 5), there will always be an action that is implementable by more than one belief. If an optimal belief implements this action, any other belief implementing the same action has to be optimal as well. This is different under continuous actions (see Section 6), where it may well be the case that the optimal belief is unique.²⁴

Notice the role of the Bayesian posterior in our model: While we call into question the mechanical updating of beliefs posited by Bayes' rule, we maintain that the Bayesian posterior, i.e., the objectively correct probability given the new information, is an important attractor in the belief choice procedure by providing the weight put onto the different states of Nature. More specifically, both the expected utility and (expected-utility) gains and losses associated with any given belief are calculated using the Bayesian posterior. Thus, our model could be interpreted as saying that people know the correct probabilities given some new piece of information, but only have the additional goal of loss avoidance in choosing their new belief. While this literal reading may be possible, we rather view our model as an "as if" representation of belief formation. People know the Bayesian posterior only at a sub-conscious level and take it sub-consciously into account when choosing their belief optimally.²⁵

²³A conventional example for the latter is a piecewise linear gain-loss function, and an example for the former is a gain-loss function that is more concave in gains than it is convex in losses or that is concave in both gains and losses.

²⁴More specifically, under continuous actions any two beliefs may well imply different actions. If different actions further induce different utility levels $u(x, a)$, the optimal belief will be unique.

²⁵Our model may therefore help "[...] to answer the question of why we all tend to hold so tightly to our own views, even when the rational part of our brains has been quite well-informed – whether through real-world experience or academic studies – that our judgment is diverging from textbook rationality" (Malmendier and Taylor, 2015, p.6).

3.2 Limiting Cases

We next discuss some general and straightforward “boundary” properties of our model. More precisely, in what follows we discuss the scenarios of no loss aversion, strong loss aversion, and certainty.

The first “polar” scenario that we consider is the case of no loss aversion, i.e., of $\lambda = 0$.

Proposition 1 (No Loss Aversion) *Suppose that the decision maker has no loss aversion, i.e., $\lambda = 0$. Then, for every $s \in S$ and $p \in Q$, it is optimal for her to choose the Bayesian posterior, i.e., $\tilde{q}(p, s) \in Q^*(p, s)$. Moreover, all other optimal beliefs must mimic $\tilde{q}(p, s)$ in the sense that $q \in Q^*(p, s)$ if and only if $A^*(q) \subseteq A^*(\tilde{q}(p, s))$.*

Proposition 1 confirms a desirable feature of our model. Intuitively, if D is not moved by loss aversion, the behavior induced by her belief choices should come as close as possible to the behavior of a standard Bayesian agent who forms her belief according to Bayes’ rule and chooses her actions so as to maximize her expected utility. Proposition 1 makes clear that this is the case: After any signal s and for any prior p , the Bayesian posterior $\tilde{q}(p, s)$ is an optimal belief for D if $\lambda = 0$. Moreover, any other optimal belief must mimic the Bayesian posterior in the sense that it triggers the same action choices (or a subset thereof).

The other polar scenario is what we refer to as “strong” loss aversion. Under strong loss aversion, D finds exactly those beliefs optimal after every signal that are action-equivalent to her prior, no matter what the latter may be.

Proposition 2 (Strong Loss Aversion) *Suppose that Assumption 1 holds. There exists a threshold level of loss aversion $\bar{\lambda} > 0$, such that for all $\lambda > \bar{\lambda}$ the decision maker exhibits “strong” loss aversion, i.e., for every $s \in S$ and $p \in Q$ we have $q \in Q^*(p, s)$ if and only if $A^*(q) = A^*(p)$.*

Thus, for sufficiently large λ D fulfills the strong loss aversion requirement. As Lemma 5 in Appendix B makes clear, all beliefs q that satisfy $A^*(q) \neq A^*(p)$ entail a loss in at least one state of the world. Hence, for a high enough λ , all beliefs satisfying $A^*(q) \neq A^*(p)$ are sub-optimal because loss avoidance becomes D’s dominant concern. Although Proposition 2 addresses an extreme scenario, it indicates how our model provides a potential explanation for *preferences for consistency*.²⁶ More specifically, if $A^*(p)$ is singleton and strong loss aversion holds, agents would choose a belief implementing the same action again, irrespective of the information received. Building on earlier work in social psychology (e.g., Cialdini, 1984), a recent experiment by Falk and Zimmermann (2015) provides clean evidence for such preferences.²⁷ Our model identifies a new reason why people want to stick to previous beliefs and actions in spite of receiving new information. This explanation provides a foundation for why agents may *desire* to act consistent, as suggested

²⁶The model of Eyster (2002) lends itself to a similar interpretation.

²⁷In their experiment, subjects entered the laboratory with an uninformative prior and received two noisy signals before they had to correctly guess the number of peas in a bowl (the action choice). While this was subjects’ first estimate in the control treatment (“CT”), subjects in the main treatment (“MT”) already had to provide an estimate after receiving the first signal without knowing that they would receive a second one. Falk and Zimmermann (2015) find that subjects in MT behave consistent with their first estimate, paying less heed to the second signal than subjects in CT. Our model can explain this pattern: Since there was no action choice in between signals’ arrival, subjects in CT should aggregate both signals before making their action-relevant belief choice. As discussed below, when priors are uninformative, agents typically follow the lead of their signal. In MT, subjects face two subsequent belief choice problems. In the first, they mainly follow the signal as priors are uninformative. But this implies that their prior is *informative* in the second belief choice problem and provokes a strong impact upon receiving the second signal.

by Falk and Zimmermann (2015). Notice, however, that despite there being a substantial fraction of subjects in their experiment behaving in line with strong loss aversion, the majority shows a milder form of consistency preferences, incorporating the signal to some extent. We will see below that our model can also account for these less extreme forms of behavior.

While the bulk of this paper restricts attention to beliefs with full support on the state space X and signals that are not fully revealing, we now make an exception in order to explore another limiting case. More specifically, we address what happens if the Bayesian posterior implies *certainty* about one particular state, e.g., because of a fully revealing signal.

Proposition 3 (Certainty) *Suppose that for some $p \in Q$ and $s \in S$ the Bayesian posterior is degenerate, i.e., there exist a state $\bar{x} \in X$ such that $\Pr(\bar{x}|p, s) = 1$. Then, the Bayesian posterior is optimal irrespective of the decision maker's loss aversion, i.e., $\tilde{q}(p, s) \in Q^*(p, s)$ for $\lambda \geq 0$.*

Proposition 3 makes clear that D 's prior being an attractor for her belief choice hinges on there being some lingering doubt about the true state of the world. It is not a goal in itself for D to match her belief to her prior. Rather, she aims to avoid expected losses from giving up the latter. This goal is traded off against her second goal, namely, to achieve maximal expected utility under the Bayesian posterior. Yet, if she is certain of a particular state \bar{x} under the Bayesian posterior, the conflict between the two goals disappears. By maximizing utility in \bar{x} , D both eliminates any loss she may feel relative to her prior and maximizes expected utility given the posterior. One way for D to do this is to adopt the Bayesian posterior itself because this causes her to choose a maximizer of $1 \times u(\bar{x}, a)$. Any other belief that is not action-equivalent to the Bayesian posterior (including the prior) is worse in terms of objective performance, and cannot be better with respect to loss aversion because *all* (objective) probability mass is on the state where the Bayesian posterior performs best.²⁸

4 Information Processing

For the remainder of the paper, we focus on the “interior” properties of our model. In order to make the analysis more tractable, we rely in Sections 4 and 5 on

Assumption 2 (Binary Setup) *The state space is $X = \{0, 1\}$ and the action space $A = \{0, 1\}$ with $u(x = 1, a = 1) = 1$, $u(x = 1, a = 0) = \alpha$, $u(x = 0, a = 0) = \beta$, and $u(x = 0, a = 1) = 0$ where $0 \leq \alpha < 1$ and $0 < \beta \leq 1$. The signal space is $S = \{0, 1\}$. Signals are informative and symmetric meaning that $\Pr(s = 0|x = 0) = \Pr(s = 1|x = 1) = \sigma$ with $\sigma \in (0.5, 1)$.*

We call $x = 1$ the *good state* and $x = 0$ the *bad state*. The state corresponds to some feature of either the external world or D 's own personality. Since the state space is binary, with little abuse of notation we can express beliefs as $f(x = 1) = f$ and $f(x = 0) = 1 - f$, where $f \in Q$. We refer to the respective events that D has received the signal $s = 1$ and $s = 0$ as the *good news* and *bad news* scenario. As for the available actions, we refer to $a = 1$ as the *bet on the good state* and to $a = 0$ as the *bet on the bad state*. In each state, the bet matching the state is preferred since $1 > \alpha$ and $\beta > 0$. Moreover, a bet that matches the state is at least as good in the good state as in the bad

²⁸Notice that Proposition 3 is a consequence of our instrumental approach. In contrast, by giving beliefs direct utility value, existing belief choice models can feature this property only as an assumption (see, e.g., Brunnermeier and Parker, 2005).

state (since $1 \geq \beta$), while a bet that does not match the state is at least as bad in the bad state as in the good state (since $0 \leq \alpha$). Strict inequalities can be driven by differences in underlying material consequences and/or state-dependent evaluation of the latter. In the limiting case where $\alpha = 0$ and $\beta = 1$, payoffs are symmetric which means that D aims to “coordinate” with Nature without favoring one state over the other. We refer to such scenarios as neutral choice contexts. Below, a key role is played by the belief $f \in Q$ for which the expected utility of betting on the good state equals that of betting on the bad state, i.e., the threshold belief that makes D indifferent between the actions in A . We denote this belief by f^A . Given Assumption 2, we have

$$0 < f^A = \frac{\beta}{1 - \alpha + \beta} < 1.$$

Several interpretations of our framework are possible. For example, $x = 1$ could be good weather and $x = 0$ bad weather. In this case, the bet on the good state consists of D taking actions appropriate for good weather (“taking no umbrella”), while the bet on the bad state amounts to preparing for bad weather (“taking an umbrella”). If the weather is good, it is better to take no umbrella ($1 > \alpha$), while the converse holds if the weather is bad ($\beta > 0$). Also, no umbrella in good weather is at least as good as an umbrella in bad weather ($1 \geq \beta$), while no umbrella in bad weather is at least as bad as an umbrella in good weather ($0 \leq \alpha$).

Secondly, “ego” may be at stake. Then, $x = 1$ could be high and $x = 0$ low ability with the bet on the good state representing a human capital investment at some monetary cost (like obtaining a university degree) and the bet on the bad state corresponding to no investment. Investing is normalized to yield a utility of one in the good state (where D is able and the investment hence worthwhile) and of zero in the bad state (where D is unable). The alternative of not investing and therefore not incurring its monetary cost yields α in the good state and β in the bad state. D prefers to invest if she is able ($1 > \alpha$) and not to invest if she is unable ($\beta > 0$). Also, going beyond Assumption 1, $\alpha \geq \beta$ seems reasonable in this context because D is likely to rank not investing and being able at least as high as not investing and being unable.²⁹

Finally, we can interpret Assumption 2 in the context of financial investments, with $x = 1$ and $x = 0$ corresponding to high and low returns of an asset, respectively. Bets on these states would be to either hold or not to hold the asset. If the asset yields high returns, it is optimal to invest in it ($\alpha < 1$), while it is better not to invest if returns are low ($\beta > 0$). Not holding the asset if returns are high is no worse than holding it if returns are low ($\alpha \geq 0$), and holding it under high returns is no worse than not holding it under low returns ($\beta \leq 1$). What we leave unspecified is whether not holding the asset under high returns is still preferable to not holding it under low returns. Considering the investment decision in isolation, $\alpha = \beta$ (i.e., equal utility from not investing in both states) seems appropriate because the immediate material consequences are the same. Alternatively, given a positive correlation between asset returns and the state of the economy (affecting other variables the investor may care about), it may well be that the former case is preferred so that $\alpha > \beta$.

²⁹There are numerous other examples where self-image concerns may arise and $\alpha \geq \beta$ seems appropriate. For instance, the state of the world could also be D’s attractiveness to another agent E and the bet on the good state (D is attractive) now consists of D approaching E, while the bet on the bad state (D is not attractive) corresponds to D abstaining from making advances to E. Then, being attractive is likely to be ranked at least as high as being unattractive, even if no attempt to conquest is made.

For some of our results, we need to characterize the belief that D adopts from her optimal beliefs $Q^*(p, s)$. Having such a point prediction facilitates our analysis in these cases.

Assumption 3 (Adopted Belief) *The decision maker's "adopted belief" $\hat{q}(p, s) \in Q^*(p, s)$ satisfies $\hat{q}(p, s) = \tilde{q}(p, s)$ if $\tilde{q}(p, s) \in Q^*(p, s)$. If $\tilde{q}(p, s) \notin Q^*(p, s)$ and $p \in Q^*(p, s)$, $\hat{q}(p, s) = p$.*

The assumption asserts that D uses a lexicographic tie-breaking rule for selecting among her optimal beliefs: If the Bayesian posterior is optimal, she adopts it. If the Bayesian posterior is not optimal, but the prior is, the latter is adopted. The case where neither Bayesian posterior nor prior is optimal does not arise below and need hence not be specified. Although Assumption 3 may not seem perfectly realistic, it minimizes deviations from the standard Bayesian benchmark since the Bayesian posterior is adopted whenever possible. Since we aim to capture non-Bayesian behavior, it is the most conservative assumption we can make in this respect.³⁰ More importantly, Assumption 3 is not crucial for our results: In Section 6, we consider a more general setup in which the action space is continuous and not binary as under Assumption 2. This yields a unique point prediction for optimal beliefs without imposing an additional assumption on adopted beliefs. In this section, we further relax the assumption of a piecewise linear gain-loss function and show that our results carry over to more complex settings that are not restricted by Assumptions 1 to 3.³¹

We now use our model to analyze information processing. We aim to account for two co-existing phenomena. Firstly, *conservatism*, which means that individuals' belief revision in response to a good or bad signal has the same sign as under Bayesian updating, but is less pronounced. Secondly, *asymmetry*, which means that good and bad news give rise to different degrees of belief revision, with updating being typically more pronounced after good news.

4.1 Conservatism

Let p be the prior probability assigned to the good state. The belief choice of D is characterized in

Lemma 1 (Optimal Beliefs) *Suppose that Assumptions 1 and 2 hold and $\lambda \geq 0$. Let*

$$p_1^D = \min \left\{ \frac{\beta(1-\sigma)(1+\lambda)}{(1-\alpha)\sigma + \beta(1-\sigma)(1+\lambda)}, f^A \right\} \text{ and } p_0^D = \max \left\{ \frac{\beta\sigma}{(1-\alpha)(1-\sigma)(1+\lambda) + \beta\sigma}, f^A \right\}.$$

If $p < p_s^D$, the decision maker prefers beliefs making her bet on the bad state, i.e., $Q^ = (0, f^A)$. If $p > p_s^D$, she prefers beliefs making her bet on the good state, i.e., $Q^* = (f^A, 1)$. And if $p = p_s^D$, $Q^* = (0, 1)$ for $\lambda < \bar{\lambda}(\sigma)$ and $Q^* = \{f^A\}$ for $\lambda > \bar{\lambda}(\sigma)$.*

Thus, p_s^D can be interpreted as D's *threshold prior* – the prior letting her choose a belief that induces indifference between the bet on the good and bad state given signal s .

Two situations can be distinguished. First, suppose that D's prior is *uninformative*, i.e., $p = f^A$. Then, if loss aversion is strong, D does not follow the lead of her signal and rather sticks to a belief inducing randomization between the two bets, i.e., her prior, since $\lambda \geq \bar{\lambda}(\sigma) \Leftrightarrow p_s^D = f^A$. On the other hand, for lower levels of loss aversion, D follows the signal because $\lambda < \bar{\lambda}(\sigma) \Leftrightarrow p_s^D \neq f^A$.

³⁰An alternative, equally conservative assumption would adopt the Bayesian posterior if optimal, and otherwise the optimal belief *closest* to the Bayesian posterior. While this alternative assumption has the advantage of delivering monotonically increasing adopted beliefs, it does not yield a closed form point prediction. We therefore prefer to impose Assumption 3, but assert that results would be similar between these two alternatives.

³¹In a sense, results become even stronger as they hold *for all* instead of *across* priors.

Hence, when priors are uninformative, agents do *not* follow the signal and deviate from Bayesian behavior only if strong loss aversion holds. Next, suppose that D's prior is *informative*, i.e., $p \neq f^A$. In this case, and just as under her prior, it is always optimal for D to choose a belief inducing a pure action strategy. Notice, however, that $p_s^D \neq p_s^B \Leftrightarrow \lambda > 0$, where p_s^B is the threshold prior at which a Bayesian agent is indifferent between betting on the good and bad state after having received signal s , and which is given by³²

$$p_1^B = \frac{\beta(1-\sigma)}{(1-\alpha)\sigma + \beta(1-\sigma)} \quad \text{and} \quad p_0^B = \frac{\beta\sigma}{(1-\alpha)(1-\sigma) + \beta\sigma}.$$

Thus, already for low levels of loss aversion there exist priors for which D chooses a belief inducing behavior that differs from a Bayesian. This property is responsible for

Proposition 4 (Conservatism) *Suppose that Assumptions 1 to 3 hold. Further, suppose that the prior distribution satisfies $\Pr(p_1^B \leq p < p_1^D) > 0$ if $s = 1$ and $\Pr(p_0^D < p \leq p_0^B) > 0$ if $s = 0$. For all $\lambda > 0$ and any signal s , average updating across priors is conservative. Moreover, $\frac{d[\|p_s^D - p_s^B\|]}{d\lambda} > 0$ for $\lambda < \bar{\lambda}(\sigma)$ and $\frac{d[\|p_s^D - p_s^B\|]}{d\lambda} = 0$ for $\lambda \geq \bar{\lambda}(\sigma)$.*

D adopts the Bayesian posterior except for priors between p_s^B and p_s^D , where she adopts her prior. Hence, with the prior distribution having positive probability in this “restriction range,” our model can account for conservatism because, after any signal, average belief revision across priors is (weakly) positive, but smaller than what is implied by Bayesian updating. The restriction range, whose size we can interpret as measuring the severity of conservatism, is increasing in D's sensitivity to loss aversion until λ has reached

$$\bar{\lambda}(\sigma) = \frac{2\sigma - 1}{1 - \sigma},$$

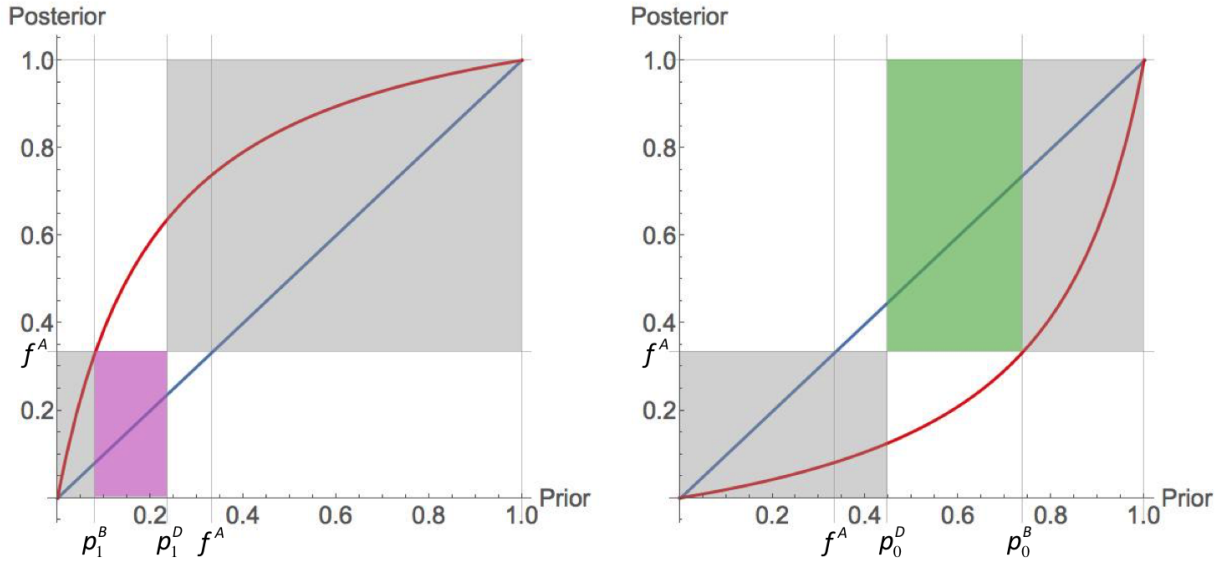
from where on it remains constant (strong loss aversion).

The intuition for Proposition 4 is straightforward. Suppose that $s = 1$. If $p < p_1^B$, a Bayesian finds it optimal to bet on the bad state, which is also what D's prior points to since we always have $p_1^B < f^A$. Clearly, D prefers betting on the bad state in this case. Since the beliefs inducing this bet include both her prior and the Bayesian posterior, Assumption 3 implies that D adopts the latter. In contrast, if $p_1^B \leq p < f^A$, a Bayesian prefers the bet on the good state, while D's prior continues to suggest the bet on the bad state. For these priors, we have the trade-off between objective performance and loss aversion that is at the heart of our model. The priors $p_1^B \leq p < p_1^D$, where $p_1^D \leq f^A$ depending on λ , are the priors where loss aversion dominates causing D to prefer the bet on the bad state. Since the beliefs inducing this bet include her prior, but no longer the Bayesian posterior, she adopts the former. Finally, if $p > p_1^D$, the force of loss aversion has become too weak or is entirely absent (if $p > f^A$). Accordingly, D finds the bet on the good state optimal, which leads her to adopt the Bayesian posterior. The intuition under bad news is analogous.

Figure 1 illustrates our results for a particular set of parameter values, given $s = 1$ (left panel) and $s = 0$ (right panel). Optimal beliefs $Q^*(p, s)$ are given by the shaded areas and the red curves represent the Bayesian posteriors as a function of p . p_s^B is pinned down by the point of intersection between the Bayesian posterior and f^A and p_s^D is the prior where optimal beliefs change.

³²The Bayesian threshold prior p_s^B is derived by applying Bayes' rule to p and setting it equal to f^A .

Figure 1: Optimal and Bayesian Beliefs



Notes: Optimal beliefs (painted areas) and Bayesian beliefs (red curves) after $s = 1$ (left panel) and $s = 0$ (right panel) as a function of the prior p given $\alpha = 0.6$, $\beta = 0.2$, $\sigma = 0.85$, and $\lambda = 2.5$.

While comparative statics with respect to loss aversion intensity are intuitive (see Proposition 4) and the same under good and bad news, the effect of varying signal precision is not. Under good news, the restriction range decreases in signal precision if the utility loss from “doing the wrong thing” is not smaller in the good than bad state, i.e., we have $\frac{d[p_1^D - p_1^B]}{d\sigma} < 0$ if $\alpha + \beta \leq 1$ and $0 < \lambda \leq \bar{\lambda}(\sigma)$. Under bad news, we have $\frac{d[p_0^B - p_0^D]}{d\sigma} < 0$ if $\alpha + \beta \geq 1$ and $0 < \lambda \leq \bar{\lambda}(\sigma)$. Thus, in neutral choice contexts (where $\alpha = 0$ and $\beta = 1$), we expect less conservatism when signals become more precise. However, when states have intrinsic value, this is not the case anymore. As we will argue below, in such settings we often expect $\alpha + \beta < 1$ to hold, implying less conservatism under more precise good news, but potentially more conservatism under more precise bad news.

What is also evident from our analysis is that the severity of conservatism not only depends on λ and σ , but also on agents’ state- and action-contingent utility levels. The dependence of the updating process on utilities raises an important question: does signal valence matter for belief updating? This is what we address next.

4.2 Asymmetry

In what follows, we use the term asymmetry to describe the phenomenon that individuals with a prior p who receive good news ($s = 1$) revise their belief more (i.e., less conservatively) than individuals with the inverse prior $1 - p$ who receive equally precise bad news ($s = 0$). Under Bayesian updating, the absolute belief changes in the two cases should be the same owing to the symmetry of priors and signals.

In our model, let

$$\Delta_1 = p_1^D - p_1^B \quad \text{and} \quad \Delta_0 = p_0^B - p_0^D$$

respectively denote the sizes of the “good news restriction range” – the width of the pink area in Figure 1 – and the “bad news restriction range” – the width of the green area. We can account

for asymmetry if and only if $\Delta_0 > \Delta_1$. To see this, note there are two cases in our model where symmetry is *not* violated: If p satisfies $p_1^B \leq p < p_1^D$ and $p_0^D < 1 - p \leq p_0^B$, symmetry is maintained because D adopts her prior both after good news given p and after bad news given the counterfactual prior $1 - p$, which means that she does not revise her belief in either case. Similarly, if p fulfills neither $p_1^B \leq p < p_1^D$ nor $p_0^D < 1 - p \leq p_0^B$, symmetry is preserved because the Bayesian posterior is adopted either way. In contrast, symmetry is violated in the other two cases: If $p_1^B \leq p < p_1^D$, but not $p_0^D < 1 - p \leq p_0^B$, D maintains p in response to good news, but adopts the Bayesian posterior in response to bad news (assuming that $1 - p$ is her prior). Likewise, if we do not have $p_1^B \leq p < p_1^D$, but $p_0^D < 1 - p \leq p_0^B$, D adopts the Bayesian posterior in response to good news, but maintains her prior $1 - p$ in response to bad news. The second type of symmetry violation occurs more often, i.e., for more priors, precisely when $\Delta_0 > \Delta_1$.³³

Proposition 5 (Asymmetry) *Suppose that Assumptions 1 to 3 hold. For all $\lambda > 0$, the size of the “bad news restriction range” exceeds (equals) the size of the “good news restriction range,” i.e., $\Delta_0 > (=) \Delta_1$, if and only if $\alpha + \beta < (=) 1$. Moreover, $\frac{d[\Delta_0 - \Delta_1]}{d\lambda} > (<) 0$ if and only if $\alpha + \beta < (>) 1$ for $\lambda < \bar{\lambda}(\sigma)$ and $\frac{d[\Delta_0 - \Delta_1]}{d\lambda} = 0$ for $\lambda \geq \bar{\lambda}(\sigma)$.*

Thus, updating is asymmetric exactly when we impose $\alpha + \beta < 1$ over and above Assumption 1. On the other hand, updating is symmetric in neutral choice contexts, where D simply wants to coordinate with the prevailing state of Nature, without ranking states in any way. In this case, $\alpha = 0$ and $\beta = 1$, which implies $\alpha + \beta = 1$. Proposition 5 also identifies when the asymmetry is reversed: whenever $\alpha + \beta > 1$, the restriction range for good news is wider than that for bad news meaning that good news are updated more conservatively on average. Thus, our model provides sharp and testable predictions on asymmetric updating. The proposition also shows that the discrepancy between the two restriction ranges – and therefore (reversed) asymmetric updating – increases in loss aversion up to $\bar{\lambda}(\sigma) = \frac{2\sigma - 1}{1 - \sigma}$, where both ranges have attained their maximal width.

The intuition for Proposition 5 is the following. The *potential* restriction range for bad news ($f^A - p_1^B$) is larger than that for good news ($p_0^B - f^A$) if and only if $\alpha + \beta < 1$ simply because $\alpha + \beta < 1$ is equivalent to $f^A < 1/2$ and by the non-linearity of the Bayesian posterior. Inside these ranges, there are two forces working against each other. On the one hand, $\alpha + \beta < 1$ is equivalent to $1 - \alpha > \beta - 0$. $1 - \alpha$ is the utility loss in the good state from betting on the bad state, whereas $\beta - 0$ is the loss in the bad state from betting on the good state. Thus, whenever $\alpha + \beta < 1$, D is exposed to a larger loss in the range for bad news because the bad news invites her to bet on the bad state, while her prior suggests the opposite, which creates scope for her feeling a loss in the good state. On the other hand, potential losses are weighted by the Bayesian posterior, and this weight tends to be larger for good news. The reason is that $f^A < 1/2$ implies relatively low priors and hence low Bayesian posteriors assigning high probability to the bad state in the potential restriction range for good news.³⁴ It can be shown, however, that the first effect always dominates the second.

As just discussed, $\alpha + \beta < 1$ means that D’s loss from betting on the wrong state is larger in the good than in the bad state. Clearly, there are cases where such a state of affairs appears less

³³Notice that the prior ranges above use $\lambda < \bar{\lambda}(\sigma)$, but we could make a similar argument using $\lambda > \bar{\lambda}(\sigma)$.

³⁴For this reason, the prior range where a loss is felt that does *not* lead to a restricted belief is also larger for bad than good news, i.e., we have $p_0^D - f^A > f^A - p_1^D \Leftrightarrow \alpha + \beta < 1$.

plausible and we may expect $\alpha + \beta > 1$ to hold. For instance, the utility loss from not carrying an umbrella in rain may seem larger than that from carrying an umbrella in sunshine. In this case, our model predicts that updating of bad news (e.g., a rainy forecast) is more pronounced than that of good news (a sunny forecast). However, the converse is likely to hold in many other settings: In a human capital context, it is likely that the monetary (and hence utility) losses from under-investing in high ability exceed those of over-investing in low ability.³⁵ In a financial investment context, $\alpha + \beta < 1$ means that the potential “upside” of an investment exceeds its “downside” (relative to the alternative of not investing). All else equal, less risk averse (or more risk seeking) individuals are therefore more likely to exhibit asymmetric updating, because for any given monetary payoff structure the upside $1 - \alpha$ of an investment is higher and the downside $\beta - 0$ is lower for them.

To illustrate, consider an investor with a utility function \tilde{u} over final wealth, who can either invest $\$w$ or not invest. If she does not invest, her final wealth will be $\$w$. If she invests, she will have either $\$w - y$ or $\$w + \varepsilon y$, where $0 < y < w$ and $\varepsilon > 0$. The investment’s upside is thus $\tilde{u}(w + \varepsilon y) - \tilde{u}(w)$ and its downside is $\tilde{u}(w) - \tilde{u}(w - y)$. Suppose first that $\varepsilon = 1$. In this case, the upside is *lower* than the downside if and only if $\tilde{u}(\cdot)$ is concave. However, for sufficiently large ε , the upside will eventually exceed the downside for any degree of risk aversion. Hence, asymmetric updating is more likely the larger (smaller) the potential gain (loss) from the investment. Notice also that the sufficiency requirement on ε is larger the more concave is $\tilde{u}(\cdot)$. So, for a given payoff structure, asymmetric updating is more likely for lower levels of risk aversion. Both these implications constitute new and testable predictions of our model. The following corollary establishes more generally how asymmetry interacts with risk attitude.

Corollary 1 (Asymmetry and Risk Attitude) *Suppose that Assumptions 1 to 3 hold. Further, suppose $v : X \times A \rightarrow \mathbb{R}$ is the final wealth position action a generates in state x . Consider two expected utility maximizers with increasing and continuous utility functions $\tilde{u}_i : \mathbb{R} \rightarrow \mathbb{R}$ and let $u_i(x, a) = \tilde{u}_i(v(x, a))$ with $i \in \{1, 2\}$. Suppose $D1$ is less risk averse (or more risk seeking) than $D2$, i.e., $\tilde{u}_2(\cdot)$ is a concave transformation of $\tilde{u}_1(\cdot)$. Then, whenever $D2$ updates bad news more conservatively than good news (so that $\Delta_0 > \Delta_1$), $D1$ will do so as well.*

Finally, we briefly discuss how our results on information processing have consequences for agents’ overconfidence. Assuming that the prior is unbiased, we say D is overconfident (underconfident) in a given situation whenever her optimal beliefs about the likelihood of the good state exceed (fall short of) the Bayesian posterior associated with the situation.³⁶ Rephrasing our previous results, under Assumptions 1 and 2, the good news restriction range induces underconfidence – shaded pink in Figure 1 – and in the bad news restriction range updating leads to overconfidence – shaded green in Figure 1. Any over- or underconfidence is “action-relevant” meaning that if p falls into one of the restriction ranges, i.e., $p_1^B \leq p < p_1^D$ or $p_0^D < p \leq p_0^B$, we have $A^*(\tilde{q}(p, s)) \neq A^*(q)$ for every $q \in Q^*(p, s)$. In terms of Figure 1, this means that the optimal

³⁵Similarly, it is likely that D feels a larger life-time utility loss from not having approached the (wo)man of his/her dreams in case of a possible conquest than from approaching him/her and being unsuccessful.

³⁶More formally, we define D as being “overconfident” (“underconfident”) after s if and only if $\inf(\sup)Q^*(p, s) > (<) \Pr(x = 1|p, s)$. Hence, if all beliefs in $Q^*(p, s)$ exceed (fall short of) the belief of a Bayesian who has the same prior and faces the same signal, D is overconfident (underconfident). This notion of overconfidence differs from one where, e.g., more than half the population believes to perform above average, and which can be consistent with Bayesian updating (Benoît and Dubra, 2011). Notice also that our notion of overconfidence directly induces overoptimism. However, the state could also refer to D ’s ability as a forecaster, and being overoptimistic about this ability may cause her to be overprecise. Hence, our notion of overconfidence may also – albeit more indirectly – be interpreted as overprecision.

beliefs are on one side of the horizontal line drawn at f^A and the Bayesian posterior on the other. Due to its instrumental nature, our model only generates this type of over- or underconfidence. A direct implication of Proposition 5 is that overconfidence is more prevalent than underconfidence whenever $\alpha + \beta < 1$. Hence, assuming a uniform prior distribution, net overconfidence (i.e., $\Delta_0 - \Delta_1$) is positive and increasing in λ if and only if $\alpha + \beta < 1$. In contrast, neutral choice contexts imply that net overconfidence is always zero. Further, by Corollary 1, D's net overconfidence is more likely to be positive the less risk averse she is, in line with the strong intuition that optimism should be positively related to risk tolerance. Finally, the location of the restriction ranges predicts overconfidence among "higher" and underconfidence among "lower" types.³⁷

4.3 Empirical Evidence

"The most common Bayesian study deals with probability estimation [...]. The primary finding has been labeled conservatism: upon receipt of new information, subjects revise their posterior probability estimates in the same direction as the optimal model, but the revision is typically too small; subjects act as if the data are less diagnostic than they truly are" (Slovic and Lichtenstein, 1971, p.693). In the psychology literature, conservatism has been established by Edwards (1968) and the bias appeared to be a robust phenomenon in experiments that followed. While Edwards (1968) attributed the bias to a human misaggregation of accurately perceived data, this literature has never settled on the exact source of the bias. Next to misaggregation, other explanations for the bias that have been suggested include misperception, anchoring, and being an experimental artefact. Loosely speaking, our model can be viewed as providing a foundation for the misaggregation and anchoring hypotheses.³⁸ However, what makes our explanation markedly distinct from existing ones is the fact that we see conservatism as a *motivated* individual trait (which allows for asymmetry) rather than the result of a mechanical updating error. Recent experimental studies in the economics literature provide evidence consistent with this characteristic.

For instance, similar to the early literature in psychology, Ambuehl and Li (2015) use a neutral choice context and ask subjects to guess the likelihood that a colored ball has been drawn from one of two urns of which subjects knew the composition. They do not only find conservatism on average, but also that responsiveness to information is individual-specific, consistent within individuals, has out-of-sample explanatory power, and is *not* caused by cognitive limitations such as one's knowledge of Bayes' rule or statistics.³⁹

Other studies allow the choice context to be non-neutral in terms of subjects' self-image. Eil and Rao (2011) investigate how subjects process noisy information about their intelligence as well as attractiveness. In both domains, they find that updating behavior conforms closely to Bayes' Rule for good news, but subjects under-update for bad news. This "good news - bad news effect" is not present in their neutral control treatment where subjects update beliefs about a randomly

³⁷Since our definition of overconfidence implicitly assumes an unbiased prior, types and priors are synonyms here. Also, notice that because $f^A < 1/2 \Leftrightarrow \alpha + \beta < 1$ higher types may well be located in the middle of the range.

³⁸Originally, anchoring was explained via insufficient adjustment away from an irrelevant anchor value (Tversky and Kahneman, 1974). While more recent literature has cast doubt on the adjustment hypothesis, Epley and Gilovich (2001, 2006) argue that insufficient adjustment does produce anchoring effects whenever anchors are *self-generated*, which is clearly what priors are.

³⁹Since subjects had an uninformative prior, on average they found only *slight* conservatism in the guessing task. But when subjects received a sequence of informative signals and therefore generally had informative priors, their updating was more conservative than predicted by their responsiveness as measured under the uninformative prior. This is consistent with our model as conservatism is least prevalent for uninformative priors (see Lemma 1).

assigned number instead of their population rank. Similarly, Möbius et al. (2014) let subjects process noisy signals about whether they scored in the top or bottom half of the population on an intelligence quiz. In a control treatment, subjects received information about a robot whose probabilistic performance matched subjects' prior belief about their own performance. The authors find that the substantial conservatism that arises when ego is at stake is reduced in the robot treatment. They also find strong evidence for good news being updated less conservatively (and thus closer to Bayes' Rule) than bad news only in their ego treatment: Subjects with a prior p who receive good news (being in the top half) revise their belief significantly more in absolute terms than subjects with the inverse prior $1 - p$ who receive bad news (being in the bottom half). When the news is instead about the robot's performance, this asymmetry disappears. In the study by Ertac (2011), subjects receive incomplete performance feedback on algebra and verbal tasks. In line with our model, the author finds that less confident subjects become pessimistic (i.e., underconfident), whereas confident ones become optimistic (i.e., overconfident) through the updating process.⁴⁰

Non-neutral choice environments can also arise due to material or other psychological consequences. Hoffman and Burks (2015) show that truckers conservatively update their beliefs about own productivity (miles driven per week) upon repeatedly receiving noisy signals about its true value. Productivity is highly consequential for earnings and beliefs are instrumental insofar as they are shown to determine workers' quitting decisions. Wiswall and Zafar (2015) find that college students revise their beliefs about future self earnings in the right direction but conservatively in response to new information about population earnings associated with different major choices.⁴¹ These belief changes have behavioral consequences in that they change subjects' major choices. Again, this emphasizes beliefs' instrumental importance, a premise that is central to our approach. Further, the authors find asymmetry as "respondents are significantly more likely to be conservative in their updating when they are informed that population earnings are lower than their prior beliefs (than when the converse is the case)" (p.3/4).⁴² Sharot et al. (2011) analyze subjects' belief updating for 80 different adverse life events – such as having a car accident, suffering a sport injury, or missing a flight – where, similar to the previous study, the news is the true empirical probability for a representative person. They report that 79% of subjects update their beliefs more in response to better-than-expected news than to worse-than-expected news.⁴³

⁴⁰While conservatism seems to be a robust finding across various contexts, the experimental evidence seems more mixed with respect to asymmetry (see, e.g., Buser et al., 2016; Coutts, 2016). In our model, conservatism occurs across all choice contexts but the presence and absence of asymmetry is affected by individual context effects. Although this may make asymmetry to appear more fragile, our results can provide guidance how to carefully test it.

⁴¹Since the signal is about public earnings but students' guesses are about private earnings, their conservatism result does not necessarily violate Bayesian updating. Here, beliefs are inelastic compared to "quasi-Bayesian" updating, which treats public signals as if they were private. Under this assumption inelastic updating can be interpreted as conservative updating.

⁴²Moreover, the authors find that belief revision is less pronounced for more experienced students. Bayesian updating can explain this result only to the extent that the same public information is less informative for more experienced students. In contrast, our model offers the additional explanation that more experienced students have more to lose from changing their beliefs (and actions): coming from an "inexperienced" situation where the utility gain from "doing the right thing" is the same in both states, it can be shown that decreasing this utility in the bad state (experience lowers β) results in more conservatism (i.e., $d\Delta_0/d\beta < 0$, provided $\beta = 1 - \alpha$ and $0 < \lambda \leq \bar{\lambda}(\sigma)$). Notice also that then $\alpha + \beta < 1$.

⁴³By identifying different regions in the brain being responsible for updating good vs. bad news, they find that the asymmetry is mainly due to a failure of the region that codes for undesirable estimation errors. Further evidence on the neurological mechanism behind asymmetric updating is, e.g., provided by Sharot et al. (2012), Moutsiana et al. (2013), or Moutsiana et al. (2015) and nicely summarized by Tali Sharot's TED talk on the "optimism bias," a.k.a. the good news - bad news effect, on https://www.ted.com/talks/tali_sharot_the_optimism_bias.

While the above evidence is in line with our *predictions*, there also exists evidence supporting our *mechanism* behind these predictions. Based on earlier neuro-scientific work, Kuhnen and Knutson (2011) hypothesize that “[...] the learning process will differ if an individual experiences a high or a low outcome, one that is better or worse than that of other investments (or actions) not chosen, or if the outcome is confirming or disconfirming of prior choices” (p.607). In their experiment, subjects received binary noisy signals about whether a risky asset was “good” or “bad” in terms of first-order stochastic dominance and then had to decide whether or not to invest in it. When the news suggests the good state, they find that conservatism is prevalent among subjects who have previously not invested, but when the news is bad, conservatism is prevalent among stock holders. Importantly, this result occurs when controlling for prior beliefs. Thus, conservatism is caused by news that contradict subjects’ prior actions rather than prior beliefs, “[...] consistent with our hypothesis that subjects choose beliefs that match their prior choices in order to avoid the negative affect caused by the admission of making a mistake” (Kuhnen and Knutson, 2011, p.620). Instead of using student data, Kuhnen et al. (2015) obtain the same results for educated, high-earning individuals and further show that the bias is mirrored in the brain.

In contrast to the direct evidence these experimental studies provide, the extant finance literature on underreaction, drift, and momentum provides rather indirect evidence on conservatism⁴⁴ and asymmetry.⁴⁵ Some of these studies (e.g., Bernard, 1993; Jegadeesh and Titman, 1993) have ruled out risk-based and other explanations and conclude that these phenomena are most likely caused by incomplete informational response. In other words, “[...] the positive autocorrelations of returns over relatively short horizons may reflect slow incorporation of news into stock prices” (Shleifer, 2000, p.112). But there is also somewhat more direct evidence: Choi and Lou (2010) find that mutual fund managers tend to boost their confidence in own stock-picking ability to a larger extent after having received confirming (public) signals than to decrease it after disconfirming ones. The bias is shown to have material consequences and leads to sub-optimal investments. Choi et al. (Forthcoming) study investors’ capital allocations and find that they move capital in the right direction but not “enough,” especially in response to managers’ *poor* performance, i.e., after bad signals. The authors attribute their findings to investors’ overweighting of priors and find it “[...] reasonable that investors exhibit an asymmetry in conservatism [...]” (p.30).

⁴⁴Conservatism causes prices – and initially investors – to underreact to news such as earnings announcements (Ball and Brown, 1968; Bernard and Thomas, 1989, 1990), share repurchase announcements (Ikenberry et al., 1995), dividend initiations/omissions (Michaely et al., 1995), stock splits (Ikenberry et al., 1996; Ikenberry and Ramnath, 2002), equity offerings (Loughran and Ritter, 1995; Spiess and Affleck-Graves, 1995), or general news relevant to fundamentals (Klibanoff et al., 1998). Choi and Hui (2014) find a similar effect regarding goals in an in-play soccer betting market. As a result of underreaction, for good (bad) news prices keep trending up (down) after the initial positive (negative) news reaction, thereby creating drift and momentum. Although momentum strategies seem to work best among stocks with low analyst coverage (Hong et al., 2000), Abarbanell and Bernard (1992) show that also analysts’ forecasts underreact to recent earnings (and that the contrary findings of DeBondt and Thaler, 1990, are not due to overreaction). Recently, Bouchaud et al. (2016) show that not only momentum but also the quality anomaly can be explained by how observed earnings forecasts of financial analysts underreact (where the measured level of “stickiness” was found to be consistent with evidence on macroeconomic forecasters).

⁴⁵Asymmetry causes a smaller initial response and larger drift, i.e., more underreaction, for bad than good news. For instance, Bernard and Thomas (1989) report that postearnings-announcement drifts are significantly more pronounced for bad than good news in small-cap stocks (primarily traded by non-institutional investors). Chan (2003) finds momentum only in stocks with public, firm-specific news, but not in those without any news. He also finds that stocks with bad news display a stronger negative drift than do stocks with good news in the positive domain. Hong et al. (2000) show that stocks with low analyst coverage react more sluggishly to bad than good news.

5 Information Preferences

In the previous section, we investigated updating behavior after receiving new information where D could not avoid obtaining the information. We now relax this assumption and ask whether or not D will acquire the information in the first place. Before potentially receiving a signal of known precision $\sigma \in (0.5, 1]$, D can now choose whether or not to expose herself to this signal. Anticipating the belief she will optimally adopt, D's expected utility, conditional on receiving signal s , is

$$EU(p, s) = \sum_x q(x)U(x, q), \quad \text{where } q \in Q^*(p, s).$$

Thus, acquiring the information yields D an expected utility of $EU^I = \sum_s \Pr(s)EU(p, s)$. If D does not demand the information, she will not receive any signal and hence cannot update. Formally, this scenario corresponds to getting an uninformative signal of precision $\sigma = 0.5$. D's expected utility in this case is $EU^{NI} = \sum_s \Pr(s)EU(p, s)$, where $p \in Q^*(p, s)$ because $\tilde{q}(p, s) = p$. The decision whether or not to acquire the information is based on a comparison between the potential cost of information and the difference in D's expected utility with and without information. Thus, if the cost $C \in \mathbb{R}$ of getting the information is sufficiently large, D will not acquire it. There is, however, always a critical \bar{C} below which (i.e., for $C \leq \bar{C}$) D will demand the information. When $\bar{C} > (<, =) 0$, D is said to be information seeking (averse, neutral).⁴⁶

5.1 Perfect Information

As a benchmark case, we will first consider a perfectly informative signal, i.e., a signal s with $\sigma = 1$. A general property of our model (see Proposition 3) is that the Bayesian posterior will be optimal irrespective of D's loss aversion if it puts all probability mass on one state.

Proposition 6 (Perfect Information Preferences) *Suppose that Assumptions 1 to 3 hold but $\sigma = 1$. For all $\lambda \geq 0$, the decision maker is strictly information seeking for all priors: If $p < f^A$, $\bar{C} = p(1 - \alpha) > 0$. And if $p > f^A$, $\bar{C} = (1 - p)\beta > 0$.*

The critical value of the information costs below which D acquires the information is *independent* of her sensitivity to losses, but dependent on her prior and the utility values $u(x, a)$. For all possible values of α , β , and p this critical value is strictly positive, indicating a positive willingness to pay for information. In other words, if noiseless information comes for free, D will acquire it. In terms of the weather example discussed earlier, if D has the chance to infer for sure whether or not it rains (e.g., by looking out of the window), she will do so and take the appropriate action. The same logic applies to non-neutral choice contexts. For instance, if an investor has the chance to get precise (and legal!) information about an asset's performance, she will always acquire it. From Proposition 6, it further follows that if the prior is below f^A , the critical value \bar{C} is increasing in p ,

⁴⁶Notice that loss aversion only influences the belief (choice of Q^*), but does not directly enter D's expected utility evaluation. In our model, only belief choice is reference-dependent, but action choices remain standard. As mentioned before, we remain standard with respect to action choice to purely isolate the effect loss aversion has on belief choice. If loss aversion additionally entered D's expected utility evaluation to determine her action regarding the information acquisition – but not concerning other actions (choice of A^*) – she would be weakly information seeking for all priors. Compared to the Bayesian benchmark, increasing λ would now increase the prior range where D was information neutral up to the point where her willingness to pay for information became zero for all priors.

but decreasing if the prior is above f^A . The reason is that at f^A the value of information is largest as D is most unsure about the appropriate action.⁴⁷

Our result for perfect information is in stark contrast to a model of belief choice featuring anticipatory utility. While Brunnermeier and Parker (2005) themselves do not cover information acquisition, Spiegel (2008) and Oster et al. (2013) have extended their model by adding a stage to the decision process in which information is chosen.⁴⁸ In doing so, Oster et al. (2013) limit their attention to perfect information and use the crucial assumption that individuals who acquire the noiseless signal cannot choose their belief anymore and simply follow the information.⁴⁹ They derive the result that agents with a large weight on anticipatory utility require compensation to obtain the information for all priors. If the weight on anticipatory utility is small, agents are information seeking for large priors, but information averse for small ones.⁵⁰ Their willingness to pay always depends on the weight on anticipatory utility. Hence, in contrast to Proposition 6, these authors find that the sign as well as the magnitude of \bar{C} depends on the importance of the psychological bias. The two models are only equivalent in the neoclassical case, i.e., if $\lambda = 0$ in our model and there is zero weight on anticipatory utility.⁵¹ In the extension by Spiegel (2008), signals are not fully revealing and hence (given binary states) never rule out any state with certainty, meaning that beliefs can be chosen no matter whether the more or less informative signal is acquired. Spiegel (2008) shows that in this case, the choice of signals is indistinguishable from that of a standard agent who is never averse to information. In what follows, we show that our model is able to generate information aversion if and only if signals are not fully revealing and in this sense yields the opposite prediction to Brunnermeier and Parker's (2005) model.

5.2 Imperfect Information

For the remainder of this section, we return to the case of $\sigma \in (0.5, 1)$. In order to isolate the effect of loss aversion on the demand for information, we first consider the case of no loss aversion.

Lemma 2 (Bayesian Information Preferences) *Suppose that Assumptions 1 to 3 hold. For $\lambda = 0$, the decision maker is weakly information seeking for all priors: If $p < p_1^B$ or $p > p_0^B$, $\bar{C} = 0$. If $p_1^B < p < f^A$, $\bar{C} = p[\beta(1 - \sigma) + \sigma(1 - \alpha)] + \beta(\sigma - 1) > 0$. And if $f^A < p < p_0^B$, $\bar{C} = p[\alpha(1 - \sigma) + \sigma(1 - \beta) - 1] + \beta\sigma > 0$.*

⁴⁷Throughout this section we exclude the special cases of indifference when separating the full prior range. This is for simplification only as including them would require tedious case distinctions that add nothing to our analysis.

⁴⁸Oster et al. (2013) use the optimal expectations framework to explain their fascinating empirical results on the behavior and beliefs of individuals at risk of Huntington disease, mainly because untested individuals are shown to remain optimistic upon receiving noisy signals (symptoms or the absence thereof) as they age. While we also strongly believe that anticipatory utility is important in such environments, it is interesting to see that new data and a corrected Bayesian posterior calculation (Oster et al., 2016) shows that this optimism finding holds true only for high motor scores (bad news). Low motor scores (good news) result in slight (albeit insignificant) pessimism. Hence, while still being inconsistent with the Bayesian approach, this new evidence seems to be in line with conservatism.

⁴⁹Agents can only choose their belief if the information is not acquired. This assumption directly roots in an assumption made by Brunnermeier and Parker (2005), stating that beliefs equal the Bayesian posterior if the latter puts all probability mass on one state. If agents could choose their belief in both cases or if the information was *slightly* imperfect (e.g., because test samples may have been mixed up), information preferences would starkly differ, i.e., agents would have a strictly positive willingness to pay for information for all priors.

⁵⁰Notice that there is a mistake in case (ii) of Oster et al.'s (2013) Proposition 3: using their notation, the critical cost value should be $(1 - p)(\delta + \Phi) - (1 - \Phi)\delta\frac{\Omega}{\Phi + \Omega}$ instead. From this correct value – but not from the incorrect one in their paper – it follows that agents always become information averse for sufficiently low priors.

⁵¹Since the “neoclassical” case of zero weight on anticipatory utility has been derived by setting $\delta = 0$ in Oster et al.'s (2013) Proposition 3, there is another mistake in case (iii) of their Proposition 5: using their notation, it should instead read $C < (1 - p)\Phi$ if $p \geq p^*$. Thus, setting $\delta = 0$ under the correct specification resembles our Proposition 6.

From Lemma 2, it follows directly that if the prior is below f^A and $\bar{C} \neq 0$, the critical cost value is increasing in p , while it is decreasing for priors above f^A . The intuition is similar to before. At priors close to f^A , the agent is most unsure whether she is taking the right decision without additional information. The only difference is that there are now “extreme” prior ranges, below p_1^B and above p_0^B , where the agent is quite confident in one state of the world, so that additional information does not have the potential to change actions. By the law of iterated expectations, the prior equals in expectation the Bayesian posterior, which makes the Bayesian agent indifferent towards information. For all other, “intermediate” priors, information is desirable.

We now turn to agents with loss aversion. Recall that if $\lambda = 0$, we have $p_s^B = p_s^D$, and D adopts the Bayesian posterior for all priors since $\Delta_s = 0$. As λ increases, p_s^D moves towards f^A until f^A is reached for $\lambda \geq \bar{\lambda}(\sigma) = \frac{2\sigma-1}{1-\sigma}$. In other words, with increasing sensitivity to losses, there are more priors for which D prefers beliefs inducing the same action as under her prior independent of the signal. If $\lambda \geq \bar{\lambda}(\sigma)$, there are no priors left for which D follows a conflicting signal, and the restriction ranges have attained their maximal width.

Proposition 7 (Imperfect Information Preferences) *Suppose that Assumptions 1 to 3 hold. For all $\lambda > 0$, the decision maker can be strictly information averse for low, but is weakly information seeking for large priors: If $p_1^B \leq p < p_1^D$, $\bar{C} = (\beta - \alpha)p(1 - p)(2\sigma - 1)$, so that $\bar{C} < (=) 0$ if and only if $\alpha > (=) \beta$. If $p_0^D < p \leq p_0^B$, $\bar{C} = p(1 - p)(2\sigma - 1) > 0$, so that $\bar{C}(\lambda > 0) > \bar{C}(\lambda = 0)$ if $\alpha \geq \beta$. For all other priors, $\bar{C}(\lambda > 0) = \bar{C}(\lambda = 0)$.*

Proposition 7 shows that within the two restriction ranges demand for information under $\lambda > 0$ is different than under $\lambda = 0$ (see Lemma 2). In the good news restriction range, i.e., for relatively low priors, D has a positive willingness to pay to *avoid* information if and only if not matching the good state still gives a higher utility than matching the bad state, which seems a plausible assumption in many non-neutral choice contexts (see our discussion following Assumption 2). The information aversion is stronger, the more D cares about the state per se rather than matching it. In contrast, in the bad news restriction range, i.e., for relatively high priors, D always has a positive willingness to pay to *acquire* information. Her willingness to pay is even larger than for a Bayesian agent if matching the good state is not worse than matching the bad state.

Intuitively, D can be information averse because she anticipates that for low priors she will update conservatively after good news, but not after bad news. Not demanding the information would lead D to stick to her prior in any case. Since getting the information has no expected utility impact under good news, it is sufficient to consider the bad news scenario. After bad news, D adopts the Bayesian posterior, which is lower than her prior, meaning that there is less probability mass on the good state than under the prior. Also, since both prior and posterior are below f^A , they both induce the bet on the bad state. Thus, whether the lower posterior or higher prior leads to higher expected utility depends on whether matching the bad state is preferred to not matching the good state. Whenever the latter is preferred, maintaining the prior is superior, and D becomes information averse (the more preferred the more averse).⁵²

On the other hand, for high priors D anticipates that her conservatism arises after bad, but not after good news, while no information makes D stick to her prior anyway. As a result, obtaining

⁵²While the same loss in expected utility after bad news (relative to acquiring no information) occurs for a Bayesian agent, there is a countervailing gain in the good news scenario (because, unlike D, a Bayesian does not restrict updating in this case), which rules out any information aversion.

information has no expected utility impact under bad news, and we can restrict attention to good news. Here, D's posterior is larger than her prior and both induce the bet on the good state. Since matching the good state always yields higher utility than not matching the bad state, the higher posterior induces a higher expected utility than the lower prior, so that D is information seeking. Compared to the Bayesian benchmark, notice that D adopts the same belief and action as a Bayesian agent under good news. However, under bad news, D's adopted belief puts a higher weight on the good state than the Bayesian posterior (because of her conservatism). In addition, while D sticks to her prior action (the bet on the good state), a Bayesian agent prefers the bet on the bad state. Therefore, a sufficient condition for D's expected utility from obtaining the information being larger than that of a Bayesian agent is that not matching the good state is still weakly preferred to matching the bad state, because this implies that after bad news the Bayesian agent, who now bets on the bad state, is losing expected utility from Bayesian updating shifting probability mass away from the good state.

Next, we analyze how agents' willingness to pay reacts when information becomes more precise.

Corollary 2 (More Precise Information) *Suppose that Assumptions 1 to 3 hold and $\lambda \geq 0$. Further, suppose that $\bar{C}(\sigma') > (<) 0$ and $\bar{C}(\sigma') \geq (<) C(\sigma') = C(\sigma'')$ for some $\sigma', \sigma'' \in (0.5, 1)$ with $\sigma'' > \sigma'$, so the decision maker is strictly information seeking (averse) and (not) willing to demand the less precise information σ' . Then, holding action choices constant, she prefers (avoids) the more precise information σ'' .*

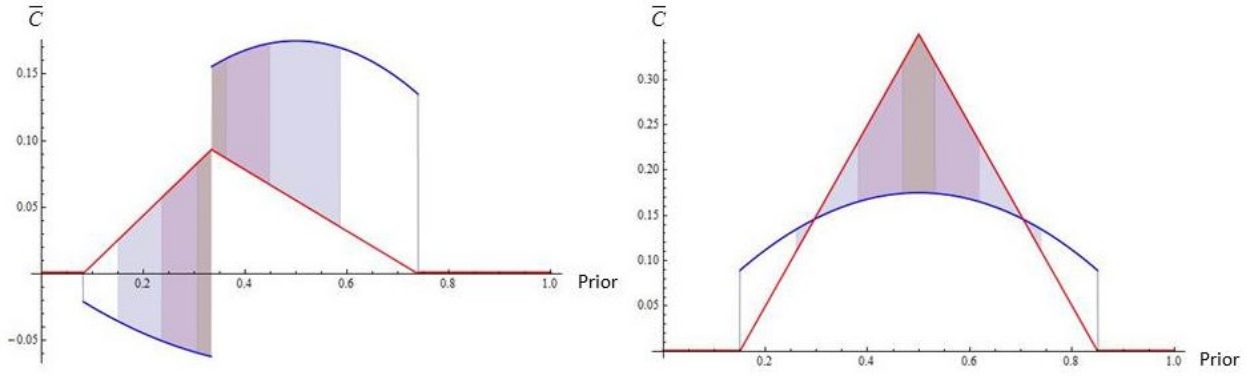
Corollary 2 makes clear that if D is information seeking and willing to acquire information, she prefers better to worse information (holding information costs constant). This result is a direct consequence of the impact σ has on \bar{C} in Lemma 2 and Proposition 7. Within all prior ranges defined above, we have $\frac{d\bar{C}}{d\sigma} > 0$ if $\bar{C} > 0$. However, since discrete jumps of \bar{C} can occur across these ranges (due to discrete changes of actions) and changing the signal precision affects their size, the corollary is restricted to priors within the same range for both σ' and σ'' , which we enforce by requiring constant action choices under σ' and σ'' .⁵³ While increasing the signal precision increases D's willingness to pay to *acquire* information when she is information seeking – so that, given that she goes, D would always want to go to the best doctor who gives diagnosis s with least noise (see Kőszegi, 2003) – the reverse happens when she is information averse. Corollary 2 further shows that in this case D's willingness to pay to *avoid* information increases for more precise signals.

Figure 2 illustrates D's willingness to pay for information (blue and red lines) as a function of her prior p . The red lines represent the case of no loss aversion, while the shaded areas indicate the extent of the departure from the Bayesian benchmark for different levels of λ (the darker the larger is λ). The left panel considers a non-neutral choice context where parameters have the same value as in Figure 1 and satisfy $\alpha > \beta$. It shows the sort of prior-dependent information aversion discussed above. The right panel considers a neutral choice context ($\alpha = 0$ and $\beta = 1$) and to which we turn next.

When matching the state is D's only concern so that $\alpha = 0$ and $\beta = 1$, Proposition 7 implies that her willingness to pay for information is weakly positive for all priors. However, D's information preferences are still distinct from a Bayesian agent:

⁵³Because such jumps do not occur without loss aversion, this condition is, however, only needed if $\lambda > 0$.

Figure 2: Willingness to Pay for Information



Notes: Willingness to pay (\bar{C}) under $\lambda = 0$ (red lines) and its departure from the Bayesian benchmark (shaded areas) induced by loss aversion (blue and red lines) as a function of the prior p . The darker the shading, the larger is loss aversion where $\lambda \in \{1, 2^{1/2}, 4\}$ or $\lambda \geq \bar{\lambda} = 42/3$. Panel A considers a non-neutral choice context where $\alpha = 0.6$, $\beta = 0.2$, and $\sigma = 0.85$. Panel B considers a neutral choice context where $\alpha = 0$, $\beta = 1$, and $\sigma = 0.85$.

Corollary 3 (Bayesian Benchmark Comparison) Suppose that Assumptions 1 to 3 hold. Let $\alpha = 0$ and $\beta = 1$. If $p_1^B \leq p < p_1^D$ ($p_0^D < p \leq p_0^B$), $\bar{C}(\lambda > 0) > \bar{C}(\lambda = 0)$ for small (large) p , but $\bar{C}(\lambda > 0) < \bar{C}(\lambda = 0)$ for large (small) p and $\lambda > \hat{\lambda}(\sigma)$. Also, $0 < (>) \frac{d\bar{C}(\lambda > 0)}{dp} < (>) \frac{d\bar{C}(\lambda = 0)}{dp} = 1$ (-1) and $0 < \frac{d\bar{C}(\lambda > 0)}{d\sigma} < \frac{d\bar{C}(\lambda = 0)}{d\sigma} = 1$.

Compared to a Bayesian agent, D overvalues information for small as well as large priors, but undervalues information for priors in the middle. The right panel of Figure 2 illustrates. As Proposition 7 already made clear that differences to Bayesian agents occur only in the two restriction ranges, Corollary 3 restricts attention to the latter (the terms “large” and “small” in Corollary 3 refer to priors’ location within these ranges). Overvaluation already occurs for small λ , but for undervaluation to occur, the restrictions ranges have to be sufficiently large. The threshold value $\hat{\lambda}(\sigma)$ is defined as the sensitivity to loss aversion such that $\bar{C}(\lambda > 0) = \bar{C}(\lambda = 0)$ at p_1^D and p_0^D , implying

$$0 < \hat{\lambda}(\sigma) = \frac{2\sigma(\sigma - 1) + \sqrt{(1 - \sigma)\sigma}}{(\sigma - 1)(-\sigma + \sqrt{(1 - \sigma)\sigma})} < \bar{\lambda}(\sigma).$$

Thus, undervaluation occurs already for values of λ smaller than the ones satisfying strong loss aversion. Another important difference to Bayesian updating, captured by Corollary 3, is that the responsiveness of D’s willingness to pay with respect to her prior confidence level is flatter (and concave instead of linear). Still, when D is most unsure she values information the most. Also, D’s willingness to pay with respect to signal precision is less responsive than for a Bayesian agent.

It is also instructive to consider our results on information preferences against the backdrop of our previous discussion on overconfidence. Suppose $C = 0$ and $\alpha > \beta$. At (low) priors that would induce underconfidence, information is avoided so that beliefs are not distorted through the updating process.⁵⁴ At (high) priors that would induce overconfidence, information is demanded allowing the updating process to in fact induce overconfidence. Thus, if agents can choose whether or not to acquire costless information, we will observe only overconfidence and no underconfidence whenever $\alpha > \beta$. Now, a rise in C can have two effects: either information is still acquired

⁵⁴The law of iterated expectations ensures that not demanding information and thereby sticking to one’s prior is expected to lead to an unbiased posterior belief: $\Pr(s = 1) \Pr(x = 1 | p, s = 1) + \Pr(s = 0) \Pr(x = 1 | p, s = 0) = p$.

and nothing changes or information is not demanded anymore. In the latter case, there is no updating and hence overconfidence cannot be present. Hence, we expect overconfidence never to increase (in the sense of affecting more priors) when information costs rise. This prediction is unaffected by the caveat that our notion of overconfidence implicitly assumes an unbiased prior in order to isolate the over- and underconfidence stemming from updating biases.⁵⁵

5.3 Empirical Evidence

The sort of prior-dependent information aversion we generate in Proposition 7 not only seems highly intuitive (Eliasz and Spiegel, 2006; Spiegel, 2008), but is also consistent with the evidence. Möbius et al. (2014) elicit subjects' willingness to pay to get feedback about their relative performance and find that a substantial fraction of subjects are willing to pay to *avoid* learning their type. Moreover, they report that more confident subjects are *causally* less information averse.⁵⁶ The same pattern is observed by Eil and Rao (2011), who also report that high types have a high willingness to pay for information while low types require a subsidy. Instead of eliciting the exact willingness to pay, Burks et al. (2013) provide subjects with the option of receiving cost-less feedback about their performance in an IQ (Raven matrices) or Numeracy test. Again, more confident subjects were more likely to demand the information. In all three experiments, ego was clearly at stake as subjects had to guess their own cognitive ability or attractiveness. As we have argued before, $\alpha > \beta$ seems a reasonable assumption in such contexts.⁵⁷ The information preferences generated by this assumption may also help explain the "Ostrich effect," which refers to the phenomenon that humans sometimes shy away from information and "bury their heads in the sand" when things look bad. Karlsson et al. (2009) test for the Ostrich effect in a finance context and find that Scandinavian as well as American investors monitor their portfolio much more frequently in rising than in falling markets. To the extent that we would expect low priors in falling, but high priors in rising markets, this observation is consistent with our model.⁵⁸

Our result on neutral choice contexts is also in line with recent evidence. Hoffman (2014) investigates in a field experiment how business professionals demand instrumental information about the price and quality of real-world websites. Since guessing correctly is these experts' only concern, the environment clearly represents the case where $\alpha = 0$ and $\beta = 1$. He finds that subjects have a strictly positive willingness to pay for information, which decreases in their confidence and the noisiness of the signal. As predicted by Corollary 3, the magnitudes of these comparative statics are less than those predicted by Bayesian updating. Moreover, Corollary 3 is consistent with the main finding of Hoffman (2014): compared to the Bayesian benchmark, subjects underpay for

⁵⁵Obviously, this is a strong assumption and due to the static nature of our model. A more realistic treatment would take into account that, over time, some signals arrive that are unavoidable whereas others may or may not be demanded. Such a dynamic treatment could potentially address how over- and underconfidence evolve from an initial uninformative prior. While this is an interesting and relevant question, it is beyond the scope of our static model.

⁵⁶In general, it will depend on the distribution of priors whether mean willingness to pay under loss aversion is positive, as it is also reported by Möbius et al. (2014). Assuming, for instance, a uniform prior distribution, one can show that mean willingness to pay is positive if $\alpha + \beta \leq 1$.

⁵⁷Although subjects in these experiments can acquire noiseless feedback, it is information about their performance on a specific task and therefore represents only a noisy signal about their true intelligence or attractiveness. Thus, it can be argued that the information is rather imperfect than perfect.

⁵⁸Notice that such information aversion may contribute to the "disposition effect," the phenomenon that investors hold on to falling assets but sell rising ones. If investors abstain from receiving information under low priors (when the economy looks bad) but not under high priors (when the economy looks good), they will not take a decision and hence hold on to their assets in bad times, but take a decision and sometimes sell their assets in good times.

valuable signals, i.e., for priors in the middle, and overpay for less valuable ones, i.e., for low and high priors. What lends additional plausibility to our model being a potential explanation of his findings is the fact that subjects, who had bought information that was contradictory to their prior guess, did not follow it one third of the time. Besides being highly significantly different from zero (the Bayesian prediction), this rate was lower for more accurate and higher for less accurate signals. Ambuehl and Li (2015) also elicit subjects' valuation for signals of varying precision in a neutral guessing task. Their results indicate two biases at work. First, similar to the finding of Hoffman (2014), subjects exhibit the "compression effect" meaning that their willingness to pay increases too little in the signal's informativeness. Second, subjects disproportionately value perfect signals as opposed to noisy ones and thus exhibit a "certainty effect." In our model, belief updating is less responsive than under Bayes' rule for noisy information, but conforms to the Bayesian benchmark for perfectly informative and perfectly uninformative signals (see Proposition 3). This property can generate a certainty effect through its impact on agents' information demand.⁵⁹

On a more fundamental level, Ambuehl and Li (2015) show that non-standard information preferences are driven by subjects' anticipation of their non-standard belief updating rather than non-standard preferences. Spiegel (2008) makes a similar suggestion: "Indeed, introspection and casual observation suggest that the phenomenon of self-deception has a lot to do with aversion to potentially unpleasant information. People who try to deceive themselves are also likely to hire 'yes men' as advisors, flip TV channels in order to avoid disturbing news, and keep 'dangerous' books out of their home. In particular, there is a strong intuition that people's choice of a biased information source over another (e.g., watching Fox News rather than BBC News) is often indicative of the kind of self-serving distortion of beliefs they are trying to attain" (p.103/104). Proposition 7 makes clear that it is *because* of their biased belief updating that individuals have biased information preferences. In particular, if an individual shies away from information, she does so because she anticipates how she would have updated her beliefs had she chosen to obtain the information. While such behavior requires information to be noisy, Propositions 3 and 6 respectively showed that neither belief updating nor information preferences are biased when information is instead fully precise. Thus, self-deception at any level, namely at the level of chosen beliefs or information demand, is present in our model if and only if the signal is noisy.^{60,61}

6 Robustness

In the two preceding sections we used three simplifying assumptions. At the cost of losing tractability, we drop them in this section to demonstrate that our previous results are robust. Assumption 1 imposed a piecewise linear gain-loss function. We now go back to the more general and possibly non-linear case. Assumption 2 imposed a binary action space. Although this as-

⁵⁹More specifically, when we compare Propositions 6 and 7 under the assumption that $\alpha = 0$ and $\beta = 1$, we find that $\lim_{\sigma \rightarrow 1} \bar{C}(\lambda > 0, \sigma < 1) = p(1-p) < p = \bar{C}(\lambda \geq 0, \sigma = 1)$ and $\lim_{\sigma \rightarrow 1} \bar{C}(\lambda > 0, \sigma < 1) = p(1-p) < 1-p = \bar{C}(\lambda \geq 0, \sigma = 1)$ in the good and bad news restriction range, respectively.

⁶⁰Sloman et al. (2010), for example, observe self-deception (about own intelligence) only when feedback on the task used vague terms rather than precise values.

⁶¹This is another noticeable difference to Brunnermeier and Parker's (2005) belief choice model, where information aversion only occurs when signals are *fully* precise, but fully precise signals are assumed to lead to unbiased beliefs. Here, agents may choose not to acquire (precise) information *because* it prohibits them – by assumption – to distort their beliefs. This is different in our model: agents may choose not to acquire (imprecise) information *because* they anticipate that they would distort their belief upon receiving it.

sumption often seems plausible, in some applications action choice may well be continuous. For instance, in the case of financial investments, instead of deciding whether or not to hold an asset, it may be appropriate to consider the choice of *how much* to invest in a risky asset. As an alternative to Assumption 2, we therefore introduce

Assumption 4 (Continuous Actions) *The state space is $X = \{0, 1\}$, the action space is $A = [0, 1]$, and the signal space is $S = \{0, 1\}$. Signals are informative and symmetric, i.e., $\Pr(s = 0 | x = 0) = \Pr(s = 1 | x = 1) = \sigma$ with $\sigma \in (0.5, 1)$. Utility in state x under action a is denoted $u(x, a)$, where $u(\cdot, a)$ is sufficiently often differentiable with (i) $u'(1, a) > 0$ and $u'(0, a) < 0$ for $a \in (0, 1)$, (ii) $u''(\cdot, a) < 0$, (iii) $\lim_{a \rightarrow x} u'(x, a) = 0$, and (iv) $u(1, 1) \geq u(0, 0)$.*

Besides making the action space continuous, Assumption 4 generalizes the utility function. Condition (i) means that utility is increasing (decreasing) in action a if $x = 1$ ($x = 0$), so that matching the state is desired. Concavity of the utility function in both states is captured by (ii), and (iii) says that marginal utility becomes zero as the action approaches the state. Jointly, these conditions assure an interior solution. Condition (iv) expresses our convention that $x = 1$ is the weakly better state. A specific utility function satisfying these conditions is

Example 1 (Quadratic Loss Function) *Utility is given by $u(1, a) = 1 - (1 - \alpha)(1 - a)^2$ and $u(0, a) = \beta - \beta(0 - a)^2$, where $0 \leq \alpha < 1$ and $0 < \beta \leq 1$.*

Using such a quadratic loss function allows us to interpret the parameters α and β as in the previous two sections. Neutral choice contexts are again captured by letting $\alpha = 0$ and $\beta = 1$. Although we derive all results in this section – formally stated and proven in Appendix B – for the more general utility function of Assumption 4, Example 1 is useful as it allows direct comparisons to our previous findings.

6.1 Information Processing

For any given belief $f \in Q$, applying Assumption 4 yields a unique optimal action which we denote by $a^*(f)$, and which is implicitly characterized by

$$-\frac{u'(1, a^*)}{u'(0, a^*)} = \frac{1 - f}{f}.$$

Notice that $\lim_{f \rightarrow 1} a^*(f) = 1$ and $\lim_{f \rightarrow 0} a^*(f) = 0$. Also, it can be shown that $a^*(f)$ is increasing in f . However, whether $a^*(f)$ is a concave, convex, or linear function – and hence whether actions are skewed towards a specific state – depends on how fast marginal utility is decreasing in each state. More specifically, with $u'''(\cdot, a) = 0$, $a^*(f)$ is concave (linear) if and only if $u''(1, a) < (=) u''(0, a)$.⁶² Thus, if marginal utility is decreasing equally fast in both states, so that states are equally good or bad, we have $a^*(f) = f$ for all f : because matching the true state is her only concern, D's optimal action follows her belief. As states become asymmetric, this property does not hold anymore: if marginal utility diminishes faster in a state, actions are skewed towards this state. Using Example 1, we have $a^* = \frac{f(1-\alpha)}{f(1-\alpha)+(1-f)\beta}$ and hence $a^* = f$ in neutral choice contexts where $\alpha = 0$ and $\beta = 1$. Actions are skewed towards the good state if and only if $\alpha + \beta < 1$ (since

⁶²Determining the shape of $a^*(f)$ requires more restrictions in case $u'''(\cdot, a) \neq 0$. We can, however, still identify that with $u'''(\cdot, a) > 0$, $a^*(f)$ is convex if $u''(1, a) > u''(0, a)$, and with $u'''(\cdot, a) < 0$, $a^*(f)$ is concave if $u''(1, a) < u''(0, a)$.

then $a^* > f$ for all f), which is reminiscent of the binary action setup (Assumption 2) where we established that $f^A < (=) 1/2$ if and only if $\alpha + \beta < (=) 1$.

However, a crucial difference to the binary action case is that for each optimal action there is now only one belief – and not a set of beliefs – that implements this action. As a result, D's optimally chosen belief yields a unique point prediction, obviating the need to impose an additional assumption, like Assumption 3, specifying D's adopted belief.

Lemma 3 (Optimal Belief) *Suppose that Assumption 4 holds and $\lambda \geq 0$. Let $q^*(p, s)$ denote the unique optimal belief given prior p and signal s , i.e., $\{q^*(p, s)\} = Q^*(p, s)$. If $\lambda < \bar{\lambda}(\sigma)$, the decision maker chooses the fixed point*

$$q^*(p, s) = \begin{cases} \frac{\Pr(x = 1|p, s)}{1 + r'(y_{p,s}^*) \Pr(x = 0|p, s)} & \text{if } s = 1 \\ \frac{(1 + r'(y_{p,s}^*)) \Pr(x = 1|p, s)}{1 + r'(y_{p,s}^*) \Pr(x = 1|p, s)} & \text{if } s = 0, \end{cases}$$

where $y_{p,s}^* = U(x \neq s, q^*(p, s)) - U(x \neq s, p)$. If $\lambda \geq \bar{\lambda}(\sigma)$, she chooses $q^*(p, s) = p$ after both signals.

The critical value of loss aversion from which on D becomes completely unresponsive to new information and sticks to her prior is the same as under binary actions and piecewise linear gain-loss utility, namely $\bar{\lambda}(\sigma) = \frac{2\sigma-1}{1-\sigma}$. No loss aversion, i.e., $\lambda = 0$, implies that $r(y) = 0$ for all y and hence $r'(y) = 0$ for all y . According to Lemma 3, D chooses $q^* = \Pr(x = 1|p, s)$ in this case. Thus, without loss aversion the Bayesian posterior is the optimally chosen belief, confirming what we already know from Proposition 1 (now the only belief that is action-equivalent to the Bayesian posterior is the Bayesian posterior itself).

If loss aversion is present ($\lambda > 0$), D's optimally chosen belief as a function of her prior has a similar shape as the Bayesian posterior, meaning it is increasing and concave (convex) for good (bad) news, but is always lower (higher) than the Bayesian benchmark. Hence, D's updating behavior goes into the same direction, but is less pronounced than under Bayes' rule. Proposition 10 in Appendix B shows that when action choice is continuous, updating is conservative for all priors. All else equal, updating becomes more conservative for larger λ : the more important loss aversion is to the decision maker, the closer is her chosen belief to her prior (until the prior is ultimately reached). For any loss aversion sensitivity, conservatism is largest for a relatively low (large) prior when $s = 1$ ($s = 0$). Thus, even though conservatism occurs for all priors, the bias is strongest for low (large) priors when the news is good (bad). In the piecewise linear case, where $r''(y) = 0$ for all y , low and large refer to below and above $1/2$, respectively (see Corollary 4 in Appendix B).

The piecewise linear case is special in another respect: As it implies that $r'(y) = \lambda$ for all $y < 0$, from Lemma 3 it can be seen that updating becomes context-independent and symmetric.⁶³ On the other hand, if $r''(y) < 0$ for all $y < 0$, Proposition 11 in Appendix B shows that updating can be asymmetric for all priors. For good news to be updated less conservatively than bad news, it is necessary (but not sufficient) that the utility loss of taking a worse action is larger in the good than in the bad state. In terms of Example 1, this requirement translates into the same condition as

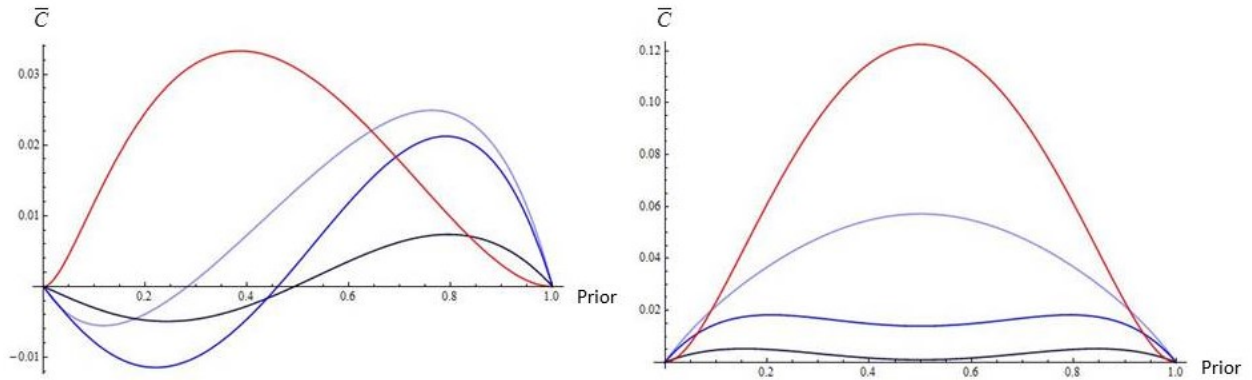
⁶³Under continuous actions, a piecewise linear gain-loss function implies that large and small losses have the same impact because optimality is determined by the marginal loss which is constant (namely λ).

under binary actions, namely $\alpha + \beta < 1$. Of course, updating can only be asymmetric if $0 < \lambda < \bar{\lambda}(\sigma)$, because the prior is chosen when $\lambda \geq \bar{\lambda}(\sigma)$ and the Bayesian posterior when $\lambda = 0$.

6.2 Information Preferences

Working in the framework of Assumption 4, also our findings on information preferences are reminiscent of what we established for the binary case. Proposition 12 in Appendix B shows that D is information seeking for all priors if the signal is fully revealing or if D is insensitive to loss aversion. For positive levels of loss aversion below $\bar{\lambda}(\sigma)$, D is information averse (neutral) for small priors if and only if not at all matching the good state is better (equal) than fully matching the bad state. In terms of Example 1, this condition translates into $\alpha > (=) \beta$. For large priors, D is always information seeking.⁶⁴ The intuition is also similar to before: For good (bad) news conservatism is maximal at rather small (large) priors (Corollary 4 in Appendix B). In terms of D's (subjective) expected utility, conservatism is a bad thing under good news, but beneficial under bad news. Taken together, this implies the distinct impact information has compared to the Bayesian benchmark.

Figure 3: Willingness to Pay for Information under Continuous Actions



Notes: Willingness to pay (\bar{C}) under $\lambda = 0$ (red lines) and under different levels of loss aversion (blue lines) as a function of the prior p . The darker a blue line, the larger is the considered loss aversion where $\lambda \in \{1, 2^{1/2}, 4\}$ and the x -axis refers to $\lambda \geq \bar{\lambda} = 4/3$. Utilities are specified using the quadratic loss function of Example 1. Panel A considers a non-neutral choice context where $\alpha = 0.6$, $\beta = 0.2$, and $\sigma = 0.85$. Panel B considers a neutral choice context where $\alpha = 0$, $\beta = 1$, and $\sigma = 0.85$.

Figure 3 illustrates D's willingness to pay for information as a function of her prior under continuous actions (Assumption 4) and piecewise linear gain-loss utility (Assumption 1), using the utility specification of Example 1 and the same parameter values as in Figures 1 and 2. The red lines represent a Bayesian agent and the blue lines displays a decision maker with different levels of loss aversion (the darker the larger is λ). The left panel considers the same non-neutral choice context as Figures 1 and 2 (satisfying $\alpha > \beta$) and shows that \bar{C} of a Bayesian agent is always positive and reaches its maximum at a point smaller than $1/2$ if and only if $\alpha + \beta < 1$. With loss aversion, D's willingness to pay is negative for low and positive for high priors, the same type of prior-dependent information aversion we already observed under binary actions. The right panel considers a neutral choice context with $\alpha = 0$ and $\beta = 1$. Again, D's willingness to pay is flatter than that of a Bayesian agent and induces overvaluation for extreme and undervaluation for less extreme priors as compared to the Bayesian benchmark.

⁶⁴For levels of loss aversion weakly above $\bar{\lambda}(\sigma)$, D is information neutral for all priors.

7 Conclusion

This paper presents a model of instrumental belief choice under loss aversion. Its core idea is that people attach value to maintaining their prior in the face of new information because they anticipate the loss (or regret) that they will feel from giving up their prior if holding on to it turns out to be optimal *ex post*. Without the influence of such loss aversion because of either insensitivity or certainty, our model reduces to the standard approach: actions generated through belief choices are then identical to those of a standard Bayesian decision maker.

With loss aversion, we identify conditions under which people exhibit a preference for consistency, meaning that their behavior becomes completely unresponsive to the arrival of new information. We also explain why and when beliefs are updated conservatively and asymmetrically. Besides closely relating our findings to empirical evidence, we provide new testable implications on the presence and absence of these biases as well as on their interaction with risk attitudes. Furthermore, we show that the anticipation of such biased updating generates information preferences which turn out to be prior-dependent: people can be information averse when their confidence in the good state is low, but are information seeking when their confidence is high. Such information preferences require a non-neutral choice context caused, for instance, by self-image concerns. In neutral choice contexts – and just as in the Bayesian model – people are always information seeking, but overvalue less valuable and undervalue more valuable information relative to the Bayesian benchmark. Despite all these differences to Bayesian updating, we demonstrate (in Appendix A) how our model is still compatible with the logistic regression model, the standard econometric approach for estimating the probability of a binary dependent variable.

At a more abstract level, the foundational idea of this paper is that the processing of new information is intrinsically context-dependent, as it depends on the specific decision problem in which new information is (likely to be) exploited. This basic tenet of our model has rich empirical implications to be explored in future work. And although it represents a departure from Bayesian decision making, by not being tailored to one specific context our model may still be considered as a “portable extension” (Rabin, 2013) to the standard approach.

A Appendix: Empirical Implementation

This appendix is concerned with the empirical implementation of our model. We first show how our approach is compatible with the (binomial) logistic regression model and thereafter propose an experimental choice task to elicit the free parameter measuring loss aversion.

A.1 Logistic Regression

It is well-known that the linear functional form of the (binomial) logistic regression model corresponds to the way Bayes' Rule uses data to update probabilities about a binary state of Nature. Thus, an important question to ask is whether our model – despite being arguably non-Bayesian – is still compatible with logistic regression. The uniqueness of the optimal belief as well as the fact that updating becomes context-independent whenever $r''(y) = 0$ for all $y < 0$ under Assumption 4 (see Lemma 3) allows us to investigate this question.

Using Bayes' Rule, the change in an *appropriate* function of beliefs does not depend on the prior belief. Of course, without such an appropriate transformation, changes in beliefs do depend on the agent's prior belief, under Bayes' Rule as well as in our model. Intuitively, when the agent is quite unsure about the true state, a signal with a given precision impacts more on her updating than in the case where the agent is quite sure about the true state of the world. So, the question is what kind of transformation makes prior beliefs irrelevant for the agent's update. Under Bayes' Rule, it is the logit transformation that achieves this. As shown by the following lemma, the logit transformation achieves the same thing in our model despite the prior's augmented importance.

Lemma 4 (Prior Invariance) *Suppose that Assumptions 1 and 4 hold and $\lambda \geq 0$. The decision maker's logit belief updating is invariant to her prior in the following sense:*

$$\text{logit}(q^*(p, s)) - \text{logit}(p) = g(\sigma, \lambda), \quad \text{where} \quad g(\cdot) = \begin{cases} \ln\left(\frac{\sigma}{(1-\sigma)(1+\lambda)}\right) & \text{if } \lambda < \bar{\lambda}(\sigma) \wedge s = 1 \\ \ln\left(\frac{(1-\sigma)(1+\lambda)}{\sigma}\right) & \text{if } \lambda < \bar{\lambda}(\sigma) \wedge s = 0 \\ 0 & \text{if } \lambda \geq \bar{\lambda}(\sigma). \end{cases}$$

Under Bayes' Rule the change of logit beliefs only depends on the precision of the signal, or more precisely, is equal to the log-likelihood ratio of the signal. Since chosen beliefs equal the Bayesian posterior without loss aversion (see Lemma 3), this property is nested in Lemma 4 by setting $\lambda = 0$. For positive values of loss aversion, Lemma 4 further shows that the change of logit beliefs additionally depends on the strength of loss aversion, or more precisely, is equal to the “ λ -adjusted” log-likelihood ratio of the signal. Importantly, however, this change is independent of the agent's prior belief. Möbius et al. (2014) provide empirical support for “prior invariance”: they report that the change of subjects' *logit* beliefs is indeed prior-independent.⁶⁵

While its empirical relevance seems noteworthy, prior invariance more fundamentally implies

⁶⁵There are two further *dynamic* properties of Bayes' Rule which Möbius et al. (2014) confirm empirically. First, priors are *sufficient statistics* for information that has arrived earlier. Second, they find that the structure of updating is *stable* over time. These two properties are fulfilled by the dynamic version of our model. Notice that Möbius et al. (2014) propose a model akin to Brunnermeier and Parker (2005) which, however, restricts potentially chosen beliefs to fulfill these three Bayesian regularities and only explores limiting infinite-horizon behavior.

Proposition 8 (Logistic Regression) *Suppose that Assumptions 1 and 4 hold and $\lambda \geq 0$. The decision maker's optimal belief $q^*(p, s)$ can be estimated using the logistic regression model, i.e.,*

$$\text{logit}(q^*(p, s)) = \beta_0 + s \times \beta_1, \quad \text{where } \beta_0 = \begin{cases} \text{logit}(p) + \ln\left(\frac{(1-\sigma)(1+\lambda)}{\sigma}\right) & \text{if } \lambda < \bar{\lambda}(\sigma) \\ \text{logit}(p) & \text{if } \lambda \geq \bar{\lambda}(\sigma) \end{cases}$$

$$\text{and } \beta_1 = \begin{cases} \ln\left(\frac{\sigma}{(1-\sigma)(1+\lambda)}\right) - \ln\left(\frac{(1-\sigma)(1+\lambda)}{\sigma}\right) & \text{if } \lambda < \bar{\lambda}(\sigma) \\ 0 & \text{if } \lambda \geq \bar{\lambda}(\sigma). \end{cases}$$

Proposition 8 shows that we can use the (binomial) logistic regression model to estimate D 's posterior belief $q^*(p, s)$, independent of whether loss aversion is present or not. The dependent variable is the posterior log odds of $x = 1$ given signal s . Under $\lambda = 0$, the conventional (Bayesian) interpretation of the coefficients applies. What changes under $\lambda > 0$ is that the interpretation of both β_0 and β_1 needs to be adjusted for D 's sensitivity to loss aversion, so that β_0 refers to the λ -adjusted log odds of $x = 1$ given $s = 0$ and β_1 to the λ -adjusted log odds ratio. However, this λ -adjustment preserves the property that the dependent variable is not only linear in s , but also linear in the parameters. Proposition 8 therefore shows that logistic regression is still applicable, even if the underlying model of belief updating does not adhere to Bayes' Rule but rather follows our approach. We next describe how to elicit λ in a separate task, so that the regression model would still have the same degrees of freedom.

A.2 Eliciting Loss Aversion

We now show how to elicit an individual's loss aversion parameter using an experimental choice task. Assume $r''(y) = 0$ for all $y < 0$ and consider the following elicitation procedure for λ . A subject D is endowed with a prior p on some state space $X = \{0, 1\}$ where $p = \Pr(x = 1) < 1/2$ (the procedure works similar for $p > 1/2$) and chooses once from the action set $A = \{0, 1\}$. Ultimately, D receives one money unit if her action choice coincides with the state, i.e., if $x = 0 \wedge a = 0$ or $x = 1 \wedge a = 1$, and no money otherwise. Since the state does not engage D 's ego, after normalization we plausibly have $u(x = 1, a = 1) = u(x = 0, a = 0) = 1$ as well as $u(x = 1, a = 0) = u(x = 0, a = 1) = 0$. This together with $p < 1/2$ implies that $a = 0$ is D 's reference action.

Using the strategy method, D is asked to make a contingent action choice before receiving a symmetric signal $s \in \{0, 1\}$ on the state. Importantly, her choice not only conditions on the signal, but also on the *signal precision*, which is not revealed to D ex ante. More specifically, D is asked to indicate the signal precision making her indifferent between the two actions after observing $s = 1$. Letting $\tilde{\sigma} \in (0.5, 1)$ be the stated precision, D is paid according to the following rule: given that $s = 1$ has materialized, action 1 (action 0) is implemented if and only if $\sigma > \tilde{\sigma}$ ($\sigma < \tilde{\sigma}$), i.e., if the actual signal precision exceeds (falls short of) the stated threshold, while the two actions are implemented with equal probability if $\sigma = \tilde{\sigma}$.

Proposition 9 (Truthful Elicitation) *In the elicitation procedure, it is a weakly dominant strategy for D to tell the truth, i.e., to choose $\tilde{\sigma} = \frac{(1-p)(1+\lambda)}{1+\lambda(1-p)} = \sigma(\lambda)$.*

To see why $\tilde{\sigma}$ pins down D's loss aversion parameter λ , we can express D's target function directly in terms of the probability of choosing action 1 denoted by π . For a given σ , D's problem is to maximize

$$V(\pi) = \frac{p\sigma}{p\sigma + (1-p)(1-\sigma)}\pi + \frac{(1-p)(1-\sigma)}{p\sigma + (1-p)(1-\sigma)}\left(1 - \pi + \lambda(1 - \pi - 1)\right)$$

where $\frac{p\sigma}{p\sigma + (1-p)(1-\sigma)}$ is the Bayesian posterior probability assigned to $x = 1$ given $s = 1$. Notice that D only feels a loss in state 0 because $p < 1/2$ implies that $a = 0$ is her reference action. We have the following indifference condition: $\frac{\partial V(\pi)}{\partial \pi} = 0 \Leftrightarrow p\sigma - (1-p)(1-\sigma)(1+\lambda) = 0$. Solving for σ yields the precision making D indifferent as $\sigma = \frac{(1-p)(1+\lambda)}{1+\lambda(1-p)}$.

B Appendix: Proofs (and Further Results)

Proof. [Proposition 1] Suppose that D holds the prior $p \in Q$ and has received the signal $s \in S$. Since $\lambda = 0$, the target function can be expressed as

$$V = \int_{a \in A^*(q)} \sum_x \Pr(x|p,s) u(x,a) dG(A^*,q) = \int_{a \in A^*(q)} \tilde{U}(a|p,s) dG(A^*,q).$$

It is immediate that $\tilde{q}(p,s) \in Q^*(p,s)$ because choosing $q = \tilde{q}(p,s)$ leads D to randomize on the actions in $A^*(\tilde{q}(p,s)) = \arg \max_{a \in A} \tilde{U}(a|p,s)$. Since $\tilde{q}(p,s)$ maximizes V by implementing randomization on $A^*(\tilde{q}(p,s))$, any $q \neq \tilde{q}(p,s)$ satisfies $q \in Q^*(p,s)$ if and only if $A^*(q) \subseteq A^*(\tilde{q}(p,s))$, i.e., if it implements randomization on a subset of $A^*(\tilde{q}(p,s))$. Hence, for every $p \in Q$ and $s \in S$, D implements only actions in $A^*(\tilde{q}(p,s))$ with positive probability regardless of which belief from $Q^*(p,s)$ she adopts. ■

Lemma 5 For any belief $q \in Q$ that satisfies $A^*(q) \neq A^*(p)$, there always exists a state of the world $x \in X$ in which D feels a loss (or regrets) abandoning p , i.e., for which we have

$$U(x,q) = \int_{a \in A^*(q)} u(x,a) dG(A^*,q) < \int_{a \in A^*(p)} u(x,a) dG(A^*,p) = U(x,p).$$

Proof. [Lemma 5] Suppose in negation that there is a belief $q \in Q$ with $A^*(q) \neq A^*(p)$ which for every state $x \in X$ satisfies $\int_{a \in A^*(q)} u(x,a) dG(A^*,q) \geq \int_{a \in A^*(p)} u(x,a) dG(A^*,p)$, which means that uniform randomization on the actions in $A^*(q)$ yields the same or higher expected utility in every state than uniform randomization on the actions in $A^*(p)$. This implies that under any belief $q \in Q$ the average expected utility of the actions in $A^*(q)$ is at least as high as the average expected utility of the actions in $A^*(p)$.

Suppose first that $A^*(q)$ is not a subset of $A^*(p)$ and consider the belief p . As just noted, the average expected utility of the actions in $A^*(q)$ under p is at least as high as the average expected utility of the actions in $A^*(p)$. By definition, the actions in $A^*(p)$ all yield the same expected utility under p , which represents the maximal level of expected utility attainable under p . Since $A^*(q)$ is not a subset of $A^*(p)$, there is at least one action in $A^*(q)$ yielding less than maximal expected utility under p . For the actions in $A^*(q)$ to yield the same or higher average utility than the actions in $A^*(p)$, this would have to be compensated by another action in $A^*(q)$ yielding more than the maximally attainable expected utility, a contradiction.

If $A^*(q)$ is a subset of $A^*(p)$ and $\int_{a \in A^*(q)} u(x,a) dG(A^*,q) > \int_{a \in A^*(p)} u(x,a) dG(A^*,p)$ for at least one $x \in X$, this implies that the average expected utility under p of the actions in $A^*(q)$ exceeds the average expected utility of the actions in $A^*(p)$. Since the actions in $A^*(p)$ attain the same, maximal expected utility under p and $A^*(q)$ is a subset of $A^*(p)$, this is a contradiction.

If $A^*(q)$ is a subset of $A^*(p)$ and $\int_{a \in A^*(q)} u(x,a) dG(A^*,q) = \int_{a \in A^*(p)} u(x,a) dG(A^*,p)$ for all $x \in X$, the expected utility under q of the actions in $A^*(q)$, which represents the maximal level of expected utility attainable under q , equals the average expected utility of the actions in $A^*(p)$. This can only be the case if either all actions that are in $A^*(p)$, but not in $A^*(q)$, attain the same expected utility under q as the actions in $A^*(q)$, a contradiction, or if at least one action that is in

$A^*(p)$, but not in $A^*(q)$, reaches a higher than maximal expected utility under q , which is also a contradiction. ■

Proof. [Proposition 2] Lemma 5 makes clear that any belief $q \in Q$ that satisfies $A^*(q) \neq A^*(p)$ entails a loss in at least one state of the world. The probability put onto this state by the Bayesian posterior $\tilde{q}(p, s)$ is guaranteed to be positive owing to our assumptions on priors (which have full support on X) and signals (which are not fully revealing).

By the scope for loss aversion assumption, at least one belief $q \in Q$ exists such that $A^*(q) \neq A^*(p)$, namely the Bayesian posterior, which satisfies $A^*(\tilde{q}(p, s)) \neq A^*(p)$ and $A^*(p) \not\subseteq A^*(\tilde{q}(p, s))$ for some $p \in Q$ and $s \in S$. Fix such a $p \in Q$ and $s \in S$. Clearly, $\tilde{q}(p, s)$ is a maximizer of the absolute part of the target function V , i.e., $\sum_x \Pr(x|p, s) U(x, q)$, which p fails to maximize because it does not induce randomization on a subset of $A^*(\tilde{q}(p, s))$. As a result, for p or any other q such that $A^*(q) = A^*(p)$ to perform better in terms of V than $\tilde{q}(p, s)$, λ must rise sufficiently above zero that avoiding the expected loss $\sum_x \Pr(x|p, s) r(U(x, \tilde{q}) - U(x, p))$ inevitably associated with $\tilde{q}(p, s)$ becomes D 's dominant concern. The same holds for any other q with $A^*(q) \neq A^*(p)$ that performs better in the absolute part of V than p .

Let $\bar{\lambda}(p, s)$ denote the threshold level of loss aversion such that for a given $p \in Q$ and $s \in S$ we have $q \in Q^*(p, s)$ if and only if $A^*(q) = A^*(p)$ for all $\lambda > \bar{\lambda}(p, s)$. This yields the threshold level $\bar{\lambda} > 0$ as $\bar{\lambda} = \max_{p \in Q, s \in S} \bar{\lambda}(p, s)$. When $\lambda > \bar{\lambda}$, "strong" loss aversion holds since D follows the lead of any p irrespective of s , thereby avoiding any loss. ■

Proof. [Proposition 3] The target function reduces to

$$V = 1 \left[\int_{a \in A^*(q)} u(\bar{x}, a) dG(A^*, q) + r \left(\int_{a \in A^*(q)} u(\bar{x}, a) dG(A^*, q) - \int_{a \in A^*(p)} u(\bar{x}, a) dG(A^*, p) \right) \right].$$

It is immediate that $\tilde{q}(p, s) \in Q^*(p, s)$ because setting $q = \tilde{q}(p, s)$ leads D to randomize on the actions in $A^*(\tilde{q}(p, s)) = \arg \max_{a \in A} 1 \times u(\bar{x}, a)$ and $r(\cdot)$ is non-decreasing. Since $\tilde{q}(p, s)$ maximizes V by only putting positive probability on actions maximizing $u(\bar{x}, a)$, any $q \neq \tilde{q}(p, s)$ satisfies $q \in Q^*(p, s)$ if and only if $A^*(q) \subseteq A^*(\tilde{q}(p, s))$, i.e., if it entails randomization on a subset of the actions maximizing $u(\bar{x}, a)$. As a result, for any $p \in Q$ and $s \in S$ such that $\Pr(\bar{x}|p, s) = 1$, D chooses only actions with positive probability that are in $A^*(\tilde{q}(p, s))$. ■

Proof. [Lemma 1] We first prove the case where $s = 1$ and distinguish between $p < p_1^D$, $p > p_1^D$, and $p = p_1^D$. If $p < p_1^D$, we have $p < f^A$ since $p_1^D \leq f^A$. Thus, the bet on the bad state is D 's reference action. For brevity, let π be the probability of betting on the good state. Given $s = 1$, D 's problem is to maximize

$$V = p\sigma \left[\pi + (1 - \pi)\alpha \right] + (1 - p)(1 - \sigma) \left[(1 - \pi)\beta + \lambda((1 - \pi)\beta - \beta) \right]. \quad (1)$$

Since $\frac{dV}{d\pi} < 0 \Leftrightarrow p < \frac{\beta(1-\sigma)(1+\lambda)}{(1-\alpha)\sigma + \beta(1-\sigma)(1+\lambda)} \geq p_1^D$, we have $\frac{dV}{d\pi} < 0$ for all $p < p_1^D$. Consequently, D prefers implementing the bet on the bad state, i.e., $Q^* = (0, f^A)$.

If $p > p_1^D$, two cases must be distinguished. Firstly, suppose $p_1^D = f^A$. We thus have $p > f^A$ and

$$V = p\sigma \left[\pi + (1 - \pi)\alpha + \lambda(\pi + (1 - \pi)\alpha - 1) \right] + (1 - p)(1 - \sigma) \left[(1 - \pi)\beta \right] \quad (2)$$

is appropriate. We have $\frac{dV}{d\pi} > 0 \Leftrightarrow p > \frac{\beta(1-\sigma)}{(1-\alpha)\sigma(1+\lambda)+\beta(1-\sigma)} < f^A$. As a result, D prefers implementing the bet on the good state, i.e., $Q^* = (f^A, 1)$. Secondly, suppose that we have $p_1^D < f^A$ implying $p_1^D = \frac{\beta(1-\sigma)(1+\lambda)}{(1-\alpha)\sigma+\beta(1-\sigma)(1+\lambda)}$. In the case where $p_1^D < p < f^A$, (1) is the appropriate target function. We have $Q^* = (f^A, 1)$ because $\frac{dV}{d\pi} > 0 \Leftrightarrow p > \frac{\beta(1-\sigma)(1+\lambda)}{(1-\alpha)\sigma+\beta(1-\sigma)(1+\lambda)} = p_1^D$. If $p > f^A$, (2) is the appropriate target function and $Q^* = (f^A, 1)$ follows from our previous discussion. Finally, if $p = f^A$, we have

$$V = \begin{cases} p\sigma \left[\pi + (1 - \pi)\alpha + \lambda(\pi + (1 - \pi)\alpha - \frac{1}{2}(1 + \alpha)) \right] + (1 - p)(1 - \sigma) \left[(1 - \pi)\beta \right] & \text{if } \pi \leq \frac{1}{2} \\ p\sigma \left[\pi + (1 - \pi)\alpha \right] + (1 - p)(1 - \sigma) \left[(1 - \pi)\beta + \lambda((1 - \pi)\beta - \frac{1}{2}\beta) \right] & \text{if } \pi > \frac{1}{2} \end{cases}$$

with V being continuous in π . In the first case where $\pi \leq \frac{1}{2}$, $\frac{dV}{d\pi} > 0$ follows from our discussion of (2). In the second case where $\pi > \frac{1}{2}$, $\frac{dV}{d\pi} > 0$ follows from our discussion of (1). We thus always have $Q^* = (f^A, 1)$ if $p > p_1^D$.

Last, if $p = p_1^D$, there are three possibilities for D's optimal beliefs. First (I), $Q^* = (0, f^A)$, suggesting the bet on the bad state. Second (II), $Q^* = (f^A, 1)$, suggesting the bet on the good state. And third (III), $Q^* = \{f^A\}$, suggesting to randomize between the bets on the good and bad state. Suppose first that $p_1^D < f^A$ and notice that $p_1^D < f^A \Leftrightarrow \lambda < \frac{2\sigma-1}{1-\sigma}$. We have $p < f^A$ since $p_1^D < f^A$. Hence, the bet on the bad state is D's reference action. D's target functions under the three possibilities above are

$$\begin{aligned} V_I &= \Pr(x = 1|p, s) \left[u(1, 0) + r(0) \right] + \Pr(x = 0|p, s) \left[u(0, 0) + r(0) \right] \\ &= \frac{\alpha p\sigma + \beta(1-p)(1-\sigma)}{p\sigma + (1-p)(1-\sigma)}, \\ V_{II} &= \Pr(x = 1|p, s) \left[u(1, 1) + r(u(1, 1) - u(1, 0)) \right] + \Pr(x = 0|p, s) \left[u(0, 1) + r(u(0, 1) - u(0, 0)) \right] \\ &= \frac{p\sigma - \beta\lambda(1-p)(1-\sigma)}{p\sigma + (1-p)(1-\sigma)}, \\ V_{III} &= \Pr(x = 1|p, s) \left[\frac{1}{2}u(1, 1) + \frac{1}{2}u(1, 0) + r\left(\frac{1}{2}u(1, 1) + \frac{1}{2}u(1, 0) - u(1, 0)\right) \right] \\ &\quad + \Pr(x = 0|p, s) \left[\frac{1}{2}u(0, 1) + \frac{1}{2}u(0, 0) + r\left(\frac{1}{2}u(0, 1) + \frac{1}{2}u(0, 0) - u(0, 0)\right) \right] \\ &= \frac{(1 + \alpha)p\sigma - \beta(\lambda - 1)(1 - p)(1 - \sigma)}{p\sigma + (1 - p)(1 - \sigma)} \times \frac{1}{2}. \end{aligned}$$

Since $p_1^D < f^A$ implies $p_1^D = \frac{\beta(1-\sigma)(1+\lambda)}{(1-\alpha)\sigma+\beta(1-\sigma)(1+\lambda)}$, we also know that $p = \frac{\beta(1-\sigma)(1+\lambda)}{(1-\alpha)\sigma+\beta(1-\sigma)(1+\lambda)}$. Plugging the latter into the three target functions yields $V_I = V_{II} = V_{III}$ and thus $Q^* = (0, 1)$. Now, suppose that $p_1^D = f^A$ and notice that $p_1^D = f^A \Leftrightarrow \lambda \geq \frac{2\sigma-1}{1-\sigma}$. In this case we have $p = f^A$ because

$p_1^D = f^A$. Therefore, D 's reference is to randomize between the bet on the good and the bet on the bad state. The target functions under the three possibilities are

$$\begin{aligned}
V_I &= \Pr(x = 1|p, s) \left[u(1, 0) + r(u(1, 0) - \frac{1}{2}u(1, 0) - \frac{1}{2}u(1, 1)) \right] \\
&\quad + \Pr(x = 0|p, s) \left[u(0, 0) + r(u(0, 0) - \frac{1}{2}u(0, 0) - \frac{1}{2}u(0, 1)) \right] \\
&= \frac{(\lambda(\alpha - 2) + 2\alpha)p\sigma + 2\beta(1 - p)(1 - \sigma)}{p\sigma + (1 - p)(1 - \sigma)} \times \frac{1}{2}, \\
V_{II} &= \Pr(x = 1|p, s) \left[u(1, 1) + r(u(1, 1) - \frac{1}{2}u(1, 0) - \frac{1}{2}u(1, 1)) \right] \\
&\quad + \Pr(x = 0|p, s) \left[u(0, 1) + r(u(0, 1) - \frac{1}{2}u(0, 0) - \frac{1}{2}u(0, 1)) \right] \\
&= \frac{2p\sigma - \beta\lambda(1 - p)(1 - \sigma)}{p\sigma + (1 - p)(1 - \sigma)} \times \frac{1}{2}, \\
V_{III} &= \Pr(x = 1|p, s) \left[\frac{1}{2}u(1, 0) + \frac{1}{2}u(1, 1) + r(0) \right] + \Pr(x = 0|p, s) \left[\frac{1}{2}u(0, 0) + \frac{1}{2}u(0, 1) + r(0) \right] \\
&= \frac{(1 + \alpha)p\sigma + \beta(1 - p)(1 - \sigma)}{p\sigma + (1 - p)(1 - \sigma)} \times \frac{1}{2}.
\end{aligned}$$

Plugging $p = f^A = \frac{\beta}{1 - \alpha + \beta}$ into the three target functions yields the following results. Given that $\lambda > \frac{2\sigma - 1}{1 - \sigma}$, we have $V_{III} > V_I$ as well as $V_{III} > V_{II}$. Thus, for $\lambda > \frac{2\sigma - 1}{1 - \sigma}$ we get $Q^* = \{f^A\}$.

We now prove the case where $s = 0$ and distinguish between $p > p_0^D$, $p < p_0^D$, and $p = p_0^D$. Consider first $p > p_0^D$. Since $p_0^D \geq f^A$, betting on the good state is D 's reference action. Taking as given $s = 0$ and letting π denote the probability of the bet on the good state, D maximizes

$$V = p(1 - \sigma) \left[\pi + (1 - \pi)\alpha + \lambda(\pi + (1 - \pi)\alpha - 1) \right] + (1 - p)\sigma \left[(1 - \pi)\beta \right]. \quad (3)$$

We have $\frac{dV}{d\pi} > 0 \Leftrightarrow p > \frac{\beta\sigma}{(1 - \alpha)(1 - \sigma)(1 + \lambda) + \beta\sigma} \leq p_0^D$. As a result, D prefers implementing the bet on the good state, i.e., $Q^* = (f^A, 1)$.

Next, consider $p < p_0^D$. Suppose first that $p_0^D = f^A$. In this case, betting on the bad state is D 's reference action, which means she maximizes

$$V = p(1 - \sigma) \left[\pi + (1 - \pi)\alpha \right] + (1 - p)\sigma \left[(1 - \pi)\beta + \lambda((1 - \pi)\beta - \beta) \right]. \quad (4)$$

We have $\frac{dV}{d\pi} < 0 \Leftrightarrow p < \frac{\beta\sigma(1 + \lambda)}{(1 - \alpha)(1 - \sigma) + \beta\sigma(1 + \lambda)} > f^A$. Thus, D prefers the bet on the bad state, i.e., $Q^* = (0, f^A)$. Suppose next that $p_0^D > f^A$ implying $p_0^D = \frac{\beta\sigma}{(1 - \alpha)(1 - \sigma)(1 + \lambda) + \beta\sigma}$. In this case, if $f^A < p < p_0^D$, (3) is the appropriate target function and we have $\frac{dV}{d\pi} < 0$ because this is equivalent to $p < \frac{\beta\sigma}{(1 - \alpha)(1 - \sigma)(1 + \lambda) + \beta\sigma} = p_0^D$. Secondly, if $p < f^A$, (4) is appropriate and $\frac{dV}{d\pi} < 0$ follows from our discussion above. Finally, consider the case $p = f^A$. We have

$$V = \begin{cases} p(1 - \sigma) \left[\pi + (1 - \pi)\alpha + \lambda(\pi + (1 - \pi)\alpha - \frac{1}{2}(1 + \alpha)) \right] + (1 - p)\sigma \left[(1 - \pi)\beta \right] & \text{if } \pi \leq \frac{1}{2} \\ p(1 - \sigma) \left[\pi + (1 - \pi)\alpha \right] + (1 - p)\sigma \left[(1 - \pi)\beta + \lambda((1 - \pi)\beta - \frac{1}{2}\beta) \right] & \text{if } \pi > \frac{1}{2} \end{cases}$$

with V being continuous in π . In the first case where $\pi \leq \frac{1}{2}$, $\frac{dV}{d\pi} < 0$ follows from our discussion of (3). In the second case where $\pi > \frac{1}{2}$, $\frac{dV}{d\pi} < 0$ follows from our discussion of (4). As a result, we always have $Q^* = (0, f^A)$ if $p < p_0^D$.

Last, consider $p = p_0^D$. There are again three possibilities for D's optimal beliefs: *I*, $Q^* = (0, f^A)$, suggesting the bet on the bad state. *II*, $Q^* = (f^A, 1)$, suggesting the bet on the good state. And *III*, $Q^* = \{f^A\}$, suggesting to randomize between the bets on the good and bad state. Suppose first that $p_0^D > f^A$ and notice that $p_0^D > f^A \Leftrightarrow \lambda < \frac{2\sigma-1}{1-\sigma}$. We have $p > f^A$ since $p_0^D > f^A$. Hence, the bet on the good state is D's reference action. D's target functions under the three possibilities above are

$$\begin{aligned}
V_I &= \Pr(x = 1|p, s) \left[u(1, 0) + r(u(1, 0) - u(1, 1)) \right] + \Pr(x = 0|p, s) \left[u(0, 0) + r(u(0, 0) - u(0, 1)) \right] \\
&= \frac{\alpha(1 + \lambda)p(1 - \sigma) + \beta(1 - p)\sigma - \lambda p(1 - \sigma)}{p(1 - \sigma) + (1 - p)\sigma}, \\
V_{II} &= \Pr(x = 1|p, s) \left[u(1, 1) + r(0) \right] + \Pr(x = 0|p, s) \left[u(0, 1) + r(0) \right] \\
&= \frac{p(1 - \sigma)}{p(1 - \sigma) + (1 - p)\sigma}, \\
V_{III} &= \Pr(x = 1|p, s) \left[\frac{1}{2}u(1, 1) + \frac{1}{2}u(1, 0) + r\left(\frac{1}{2}u(1, 1) + \frac{1}{2}u(1, 0) - u(1, 1)\right) \right] \\
&\quad + \Pr(x = 0|p, s) \left[\frac{1}{2}u(0, 1) + \frac{1}{2}u(0, 0) + r\left(\frac{1}{2}u(0, 1) + \frac{1}{2}u(0, 0) - u(0, 1)\right) \right] \\
&= \frac{(\alpha(2 + \lambda)(1 - \lambda) + \sigma(\lambda - 1 - \beta))p + \beta\sigma}{p(1 - \sigma) + (1 - p)\sigma} \times \frac{1}{2}.
\end{aligned}$$

Since $p_0^D > f^A$ implies $p_0^D = \frac{\beta\sigma}{(1-\alpha)(1-\sigma)(1+\lambda)+\beta\sigma}$, we also know that $p = \frac{\beta\sigma}{(1-\alpha)(1-\sigma)(1+\lambda)+\beta\sigma}$. Plugging the latter into the three target functions yields $V_I = V_{II} = V_{III}$ and thus $Q^* = (0, 1)$. Now, suppose that $p_0^D = f^A$ and notice that $p_0^D = f^A \Leftrightarrow \lambda \geq \frac{2\sigma-1}{1-\sigma}$. In this case we have $p = f^A$ because $p_0^D = f^A$. Therefore, D's reference is to randomize between the bet on the good and the bet on the bad state. The target functions under the three possibilities are

$$\begin{aligned}
V_I &= \Pr(x = 1|p, s) \left[u(1, 0) + r\left(u(1, 0) - \frac{1}{2}u(1, 0) - \frac{1}{2}u(1, 1)\right) \right] \\
&\quad + \Pr(x = 0|p, s) \left[u(0, 0) + r\left(u(0, 0) - \frac{1}{2}u(0, 0) - \frac{1}{2}u(0, 1)\right) \right] \\
&= \frac{\alpha(2 + \lambda)p(1 - \sigma) + 2(\beta(1 - p)\sigma - \lambda p(1 - \sigma))}{p(1 - \sigma) + (1 - p)\sigma} \times \frac{1}{2}, \\
V_{II} &= \Pr(x = 1|p, s) \left[u(1, 1) + r\left(u(1, 1) - \frac{1}{2}u(1, 0) - \frac{1}{2}u(1, 1)\right) \right] \\
&\quad + \Pr(x = 0|p, s) \left[u(0, 1) + r\left(u(0, 1) - \frac{1}{2}u(0, 0) - \frac{1}{2}u(0, 1)\right) \right] \\
&= \frac{\beta\lambda\sigma - p(2 + (\beta\lambda - 2)\sigma)}{(1 - p)(1 - \sigma)} \times \frac{1}{2}, \\
V_{III} &= \Pr(x = 1|p, s) \left[\frac{1}{2}u(1, 0) + \frac{1}{2}u(1, 1) + r(0) \right] + \Pr(x = 0|p, s) \left[\frac{1}{2}u(0, 0) + \frac{1}{2}u(0, 1) + r(0) \right] \\
&= \frac{\beta\sigma + p(1 + \alpha(1 - \sigma) - \sigma(1 + \beta))}{p(1 - \sigma) + (1 - p)\sigma} \times \frac{1}{2}.
\end{aligned}$$

Plugging $p = f^A = \frac{\beta}{1-\alpha+\beta}$ into the three target functions yields the following results. Given that $\lambda > \frac{2\sigma-1}{1-\sigma}$, we have $V_{III} > V_I$ as well as $V_{III} > V_{II}$. Thus, for $\lambda > \frac{2\sigma-1}{1-\sigma}$ we get $Q^* = \{f^A\}$. ■

Proof. [Proposition 4] We apply Assumption 3 to Lemma 1 in order to show what D's adopted beliefs are. We start with the case where $s = 1$ and distinguish three cases: $p < p_1^D$, $p > p_1^D$, and $p = p_1^D$. Suppose first that $p < p_1^D$. Since $p_1^D > (=) p_1^B \Leftrightarrow \lambda > (=) 0$, there are two possibilities. When $p < p_1^B$, we have $0 < \Pr(x = 1|p, s) < f^A$ which implies $\hat{q} = \Pr(x = 1|p, s)$ by Assumption 2. In contrast, if $p_1^B \leq p < p_1^D$, we have $\Pr(x = 1|p, s) \geq f^A$. By Assumption 2, we now have $\hat{q} = p$ since the posterior is no longer part of D's preferred beliefs. Next, suppose that $p > p_1^D$. Since $p_1^D > (=) p_1^B \Leftrightarrow \lambda > (=) 0$, $f^A < \Pr(x = 1|p, s) < 1$ holds for these priors and we therefore have $\hat{q} = \Pr(x = 1|p, s)$ by Assumption 3. Last, suppose that $p = p_1^D$. Since $p_1^D > (=) p_1^B \Leftrightarrow \lambda > (=) 0$, $f^A < \Pr(x = 1|p, s) < 1$ holds for these priors. By Assumption 3, we thus have $\hat{q} = \Pr(x = 1|p, s)$ for $\lambda < \frac{2\sigma-1}{1-\sigma}$, but $\hat{q} = p$ for $\lambda > \frac{2\sigma-1}{1-\sigma}$ because the Bayesian posterior is no longer part of D's preferred beliefs.

We now turn to the case where $s = 0$ and again distinguish the following cases: $p < p_0^D$, $p > p_0^D$, and $p = p_0^D$. First, when $p < p_0^D$, $0 < \Pr(x = 1|p, s) < f^A$ holds for these priors because $p_0^D < (=) p_0^B \Leftrightarrow \lambda > (=) 0$. As a result, Assumption 3 implies $\hat{q} = \Pr(x = 1|p, s)$. Second, suppose that $p > p_0^D$. Two possibilities emerge as $p_0^D < (=) p_0^B \Leftrightarrow \lambda > (=) 0$. When $p > p_0^B$, we have $f^A < \Pr(x = 1|p, s) < 1$ which implies $\hat{q} = \Pr(x = 1|p, s)$ by Assumption 2. In contrast, if $p_0^D < p \leq p_0^B$, we have $\Pr(x = 1|p, s) \leq f^A$. By Assumption 2, we now have $\hat{q} = p$ since the posterior is no longer part of D's preferred beliefs. Third, suppose that $p = p_0^D$. Since $p_0^D < (=) p_0^B \Leftrightarrow \lambda > (=) 0$, $0 < \Pr(x = 1|p, s) < f^A$ holds for these priors. By Assumption 3, we thus have $\hat{q} = \Pr(x = 1|p, s)$ for $\lambda < \frac{2\sigma-1}{1-\sigma}$, but $\hat{q} = p$ for $\lambda > \frac{2\sigma-1}{1-\sigma}$ as the Bayesian posterior is no longer part of D's preferred beliefs.

To summarize, when

- (i) $s = 1$: $\hat{q} = p$ if $p_1^B \leq p < (\leq) p_1^D$ and $\hat{q} = \Pr(x = 1|p, s)$ if $p < p_1^B$ or $p \geq (>) p_1^D$ for $\lambda < (>) \bar{\lambda}(\sigma)$;
- (ii) $s = 0$: $\hat{q} = p$ if $p_0^D < (\leq) p \leq p_0^B$ and $\hat{q} = \Pr(x = 1|p, s)$ if $p \leq (<) p_0^D$ or $p > p_0^B$ for $\lambda < (>) \bar{\lambda}(\sigma)$.

Thus, when $s = 1$, for all $\lambda > 0$ D adopts her prior for all $p \in [p_1^B, p_1^D)$ and this range is strictly positive since $p_1^D > (=) p_1^B \Leftrightarrow \lambda > (=) 0$. Hence, $\Pr(p \in [p_1^B, p_1^D)) > 0$ is sufficient to ensure that D is conservative when we average over the prior distribution. Similarly, when $s = 0$, for all $\lambda > 0$ D adopts her prior for all $p \in (p_0^D, p_0^B]$ and this range is strictly positive since $p_0^D < (=) p_0^B \Leftrightarrow \lambda > (=) 0$. Thus, when we average over the prior distribution, $\Pr(p \in (p_0^D, p_0^B]) > 0$ is sufficient to ensure that D is conservative.

The last part of Proposition 4 follows from observing that p_s^B as well as f^A are independent of λ , that $p_s^D = f^A \Leftrightarrow \lambda \geq \frac{2\sigma-1}{1-\sigma}$, and that

$$\begin{aligned} \frac{d}{d\lambda} \left[\frac{\beta(1-\sigma)(1+\lambda)}{(1-\alpha)\sigma + \beta(1-\sigma)(1+\lambda)} \right] &= \frac{(1-\alpha)\beta(1-\sigma)\sigma}{[\beta(1+\lambda)(\sigma-1) + (\alpha-1)\sigma]^2} > 0 \quad \text{and} \\ \frac{d}{d\lambda} \left[\frac{\beta\sigma}{(1-\alpha)(1-\sigma)(1+\lambda) + \beta\sigma} \right] &= \frac{-(1-\alpha)\beta(1-\sigma)\sigma}{[(\alpha-1)(1+\lambda)(\sigma-1) + \beta\sigma]^2} < 0 \end{aligned}$$

where both inequalities follow from Assumption 2. ■

Proof. [Proposition 5] Notice first that $\Delta_0 = \Delta_1 = 0$ if $\lambda = 0$. Now, suppose that $0 < \lambda < \frac{2\sigma-1}{1-\sigma}$. In this case, comparing Δ_0 and Δ_1 yields

$$\begin{aligned}
\Delta_0 > \Delta_1 &\Leftrightarrow \frac{\beta\sigma}{(1-\alpha)(1-\sigma) + \beta\sigma} - \frac{\beta\sigma}{(1-\alpha)(1-\sigma)(1+\lambda) + \beta\sigma} \\
&> \frac{\beta(1-\sigma)(1+\lambda)}{(1-\alpha)\sigma + \beta(1-\sigma)(1+\lambda)} - \frac{\beta(1-\sigma)}{(1-\alpha)\sigma + \beta(1-\sigma)} \\
&\Leftrightarrow \frac{(\alpha-1)\beta\lambda(\sigma-1)\sigma}{(1+\alpha(\sigma-1) + (\beta-1)\sigma)(1+\lambda + \alpha(1+\lambda)(\sigma-1) - \sigma + \beta\sigma - \lambda\sigma)} \\
&> \frac{(\alpha-1)\beta\lambda(\sigma-1)\sigma}{(\beta(\sigma-1) + (\alpha-1)\sigma)\beta(1+\lambda(\sigma-1) + (\alpha-1)\sigma)} \\
&\Leftrightarrow (1+\alpha(\sigma-1) + (\beta-1)\sigma)(1+\lambda + \alpha(1+\lambda)(\sigma-1) - \sigma + \beta\sigma - \lambda\sigma) \\
&\quad - (\beta(\sigma-1) + (\alpha-1)\sigma)\beta(1+\lambda(\sigma-1) + (\alpha-1)\sigma) < 0 \\
&\Leftrightarrow (1-2\alpha + \alpha^2 - \beta^2)(1+\lambda(\sigma-1)^2 - 2\sigma) < 0.
\end{aligned}$$

We have $1 + \lambda(\sigma-1)^2 - 2\sigma < 0$ because it is equivalent to $\lambda < \frac{2\sigma-1}{(1-\sigma)^2}$ and since $\frac{2\sigma-1}{(1-\sigma)^2} > \frac{2\sigma-1}{1-\sigma}$ (as $\sigma \in (0.5, 1)$ by Assumption 2). Moreover, $1 - 2\alpha + \alpha^2 - \beta^2 > (=) 0 \Leftrightarrow \alpha + \beta < (=) 1$. Now, suppose that $\lambda \geq \frac{2\sigma-1}{1-\sigma}$, so that

$$\begin{aligned}
\Delta_0 > \Delta_1 &\Leftrightarrow \frac{\beta\sigma}{(1-\alpha)(1-\sigma) + \beta\sigma} - \frac{\beta}{1-\alpha+\beta} > \frac{\beta}{1-\alpha+\beta} - \frac{\beta(1-\sigma)}{(1-\alpha)\sigma + \beta(1-\sigma)} \\
&\Leftrightarrow \frac{(1-\alpha)\beta(2\sigma-1)}{(1-\alpha+\beta)(1+\alpha(\sigma-1) + \sigma(\beta-1))} > \frac{(1-\alpha)\beta(2\sigma-1)}{(\alpha-\beta-1)(\sigma(\alpha-1) + \beta(\sigma-1))} \\
&\Leftrightarrow (1-\alpha+\beta)(1+\alpha(\sigma-1) + \sigma(\beta-1)) - (\alpha-\beta-1)(\sigma(\alpha-1) + \beta(\sigma-1)) < 0 \\
&\Leftrightarrow -(1-2\alpha + \alpha^2 - \beta^2)(2\sigma-1) < 0.
\end{aligned}$$

We have $2\sigma-1 > 0$ (as $\sigma \in (0.5, 1)$ by Assumption 2) and $1 - 2\alpha + \alpha^2 - \beta^2 > (=) 0 \Leftrightarrow \alpha + \beta < (=) 1$.

Because $p_s^D = f^A \Leftrightarrow \lambda \geq \frac{2\sigma-1}{1-\sigma}$, any further increases of λ do not affect the width of the restriction ranges as not only f^A but also p_s^B is independent of λ . Consequently, $\frac{d[\Delta_0 - \Delta_1]}{d\lambda} = 0$ if $\lambda \geq \frac{2\sigma-1}{1-\sigma}$. If $\lambda < \frac{2\sigma-1}{1-\sigma}$, we get

$$\begin{aligned}
\frac{d[\Delta_0 - \Delta_1]}{d\lambda} &= \frac{(1-\alpha)\beta(1-\sigma)\sigma}{[(1-\alpha)(1+\lambda)(1-\sigma) + \beta\sigma]^2} - \frac{(1-\alpha)\beta(1-\sigma)\sigma}{[\beta(1+\lambda)(1-\sigma) + (1-\alpha)\sigma]^2} > 0 \\
&\Leftrightarrow [(1-\alpha)(1+\lambda)(1-\sigma) + \beta\sigma]^2 < [\beta(1+\lambda)(1-\sigma) + (1-\alpha)\sigma]^2 \\
&\Leftrightarrow (\alpha + \beta - 1)(2\sigma - 1 + \lambda(\sigma - 1)) < 0,
\end{aligned}$$

where the last equivalence follows because $(1-\alpha)(1+\lambda)(1-\sigma) + \beta\sigma > 0$ as well as $\beta(1+\lambda)(1-\sigma) + (1-\alpha)\sigma > 0$ by Assumption 2. We have $2\sigma - 1 + \lambda(\sigma - 1) > 0$ because it is equivalent to $\lambda < \frac{2\sigma-1}{1-\sigma}$. Moreover, it is immediate that $\alpha + \beta - 1 < (>) 0 \Leftrightarrow \alpha + \beta < (>) 1$. ■

Proof. [Corollary 1] Notice first that by Assumption 2, because $u(\cdot)$ is strictly increasing, and since $u_i(x, a) = \tilde{u}_i(v(x, a))$, we have $v(1, 1) > v(1, 0) \geq v(0, 1)$ and $v(1, 1) \geq v(0, 0) > v(0, 1)$. Normalizing as in Assumption 2, D_i (with $i \in \{1, 2\}$) therefore assigns $u_i(1, 1) = 1$, $u_i(1, 0) = \alpha_i$, $u_i(0, 0) = \beta_i$, and $u_i(0, 1) = 0$. Since, $\tilde{u}_2(\cdot)$ is more concave than $\tilde{u}_1(\cdot)$, we have $\alpha_2 \geq \alpha_1$ and

$\beta_2 \geq \beta_1$. Thus, whenever D2 updates asymmetrically, D1 will do so as well, because $1 - \alpha_2 > \beta_2 \Rightarrow 1 - \alpha_1 > \beta_1$. ■

Proof. [Proposition 6] Since $\sigma = 1$, we have $\Pr(x = s|p, s) = 1$. By Proposition 3, we thus know that Assumption 3 implies $\hat{q} = \Pr(x = 1|p, s) = 0 \Leftrightarrow s = 0$ and $\hat{q} = \Pr(x = 1|p, s) = 1 \Leftrightarrow s = 1$ for all p and for all $\lambda \geq 0$. Hence, we have $a = s$ when D acquires information. But when D does not demand information, we distinguish two cases. First, if $p < f^A$ D's prior suggests the bet on the bad state and, second, if $p > f^A$ the prior suggests the bet on the good state. In the first case D demands information if and only if

$$\begin{aligned} & \Pr(s = 1)u(1, 1) + \Pr(s = 0)u(0, 0) - C \geq pu(1, 0) + (1 - p)u(0, 0) \\ \Leftrightarrow & p + (1 - p)\beta - C \geq p\alpha + (1 - p)\beta \\ \Leftrightarrow & C \leq p(1 - \alpha), \end{aligned}$$

where the right-hand side (RHS) is positive since $p > 0$ and by Assumption 2. In the second case, D demands information if and only if

$$\begin{aligned} & \Pr(s = 1)u(1, 1) + \Pr(s = 0)u(0, 0) - C \geq pu(1, 1) + (1 - p)u(0, 1) \\ \Leftrightarrow & p + (1 - p)\beta - C \geq p \\ \Leftrightarrow & C \leq (1 - p)\beta, \end{aligned}$$

where the RHS is positive since $p < 1$ and by Assumption 2. ■

Proof. [Lemma 2] With $\lambda = 0$, the Bayesian posterior is always in the set of optimal beliefs (see Proposition 1) and thus always adopted (by Assumption 3). We therefore always have $\hat{q} = \Pr(x = 1|p, s)$. Since $\sigma \in (0.5, 1)$, we need to distinguish four cases. First, when $p < p_1^B$ D's prior and the Bayesian posterior suggest the bet on the bad state. Second, when $p > p_0^B$ both prior and Bayesian posterior suggest the bet on the good state. Third, when $p_1^B < p < f^A$, the prior suggests the bet on the bad state while the Bayesian posterior suggests $a = s$. Fourth, when $f^A < p < p_0^B$, D's prior suggests the bet on the good state while the Bayesian posterior suggests $a = s$. In the first case, D demands information if and only if

$$\begin{aligned} & \Pr(s = 1)[\Pr(x = 1|p, s = 1)u(1, 0) + \Pr(x = 0|p, s = 1)u(0, 0)] \\ & + \Pr(s = 0)[\Pr(x = 1|p, s = 0)u(1, 0) + \Pr(x = 0|p, s = 0)u(0, 0)] - C \geq pu(1, 0) + (1 - p)u(0, 0) \\ \Leftrightarrow & [p\sigma + (1 - p)(1 - \sigma)]\left[\frac{p\sigma}{p\sigma + (1 - p)(1 - \sigma)}\alpha + \frac{(1 - p)(1 - \sigma)}{p\sigma + (1 - p)(1 - \sigma)}\beta\right] \\ & + [(1 - p)\sigma + p(1 - \sigma)]\left[\frac{p(1 - \sigma)}{p(1 - \sigma) + (1 - p)\sigma}\alpha + \frac{(1 - p)\sigma}{p(1 - \sigma) + (1 - p)\sigma}\beta\right] - C \geq p\alpha + (1 - p)\beta \\ \Leftrightarrow & C \leq 0. \end{aligned}$$

Similarly, in the second case, D demands information if and only if

$$\begin{aligned}
& \Pr(s = 1)[\Pr(x = 1|p, s = 1)u(1, 1) + \Pr(x = 0|p, s = 1)u(0, 1)] \\
& + \Pr(s = 0)[\Pr(x = 1|p, s = 0)u(1, 1) + \Pr(x = 0|p, s = 0)u(0, 1)] - C \geq pu(1, 1) + (1 - p)u(0, 1) \\
\Leftrightarrow & [p\sigma + (1 - p)(1 - \sigma)]\left[\frac{p\sigma}{p\sigma + (1 - p)(1 - \sigma)}\right] + [(1 - p)\sigma + p(1 - \sigma)]\left[\frac{p(1 - \sigma)}{p(1 - \sigma) + (1 - p)\sigma}\right] - C \geq p \\
\Leftrightarrow & C \leq 0.
\end{aligned}$$

In the third case, D demands information if and only if

$$\begin{aligned}
& \Pr(s = 1)[\Pr(x = 1|p, s = 1)u(1, 1) + \Pr(x = 0|p, s = 1)u(0, 1)] \\
& + \Pr(s = 0)[\Pr(x = 1|p, s = 0)u(1, 0) + \Pr(x = 0|p, s = 0)u(0, 0)] - C \geq pu(1, 0) + (1 - p)u(0, 0) \\
\Leftrightarrow & [p\sigma + (1 - p)(1 - \sigma)]\left[\frac{p\sigma}{p\sigma + (1 - p)(1 - \sigma)}\right] \\
& + [(1 - p)\sigma + p(1 - \sigma)]\left[\frac{p(1 - \sigma)}{p(1 - \sigma) + (1 - p)\sigma}\alpha + \frac{(1 - p)\sigma}{p(1 - \sigma) + (1 - p)\sigma}\beta\right] - C \geq p\alpha + (1 - p)\beta \\
\Leftrightarrow & C \leq p[\beta(1 - \sigma) + \sigma(1 - \alpha)] + \beta(\sigma - 1),
\end{aligned}$$

where the RHS is positive since $p > p_1^B$. In the fourth case, D demands information if and only if

$$\begin{aligned}
& \Pr(s = 1)[\Pr(x = 1|p, s = 1)u(1, 1) + \Pr(x = 0|p, s = 1)u(0, 1)] \\
& + \Pr(s = 0)[\Pr(x = 1|p, s = 0)u(1, 0) + \Pr(x = 0|p, s = 0)u(0, 0)] - C \geq pu(1, 1) + (1 - p)u(0, 1) \\
\Leftrightarrow & [p\sigma + (1 - p)(1 - \sigma)]\left[\frac{p\sigma}{p\sigma + (1 - p)(1 - \sigma)}\right] \\
& + [(1 - p)\sigma + p(1 - \sigma)]\left[\frac{p(1 - \sigma)}{p(1 - \sigma) + (1 - p)\sigma}\alpha + \frac{(1 - p)\sigma}{p(1 - \sigma) + (1 - p)\sigma}\beta\right] - C \geq p \\
\Leftrightarrow & C \leq p[\alpha(1 - \sigma) + \sigma(1 - \beta) - 1] + \beta\sigma,
\end{aligned}$$

where the RHS is positive since $p < p_0^B$. ■

Proof. [Proposition 7] Since D's prior and the Bayesian posterior suggest the same bet when $p < p_1^B$ or $p > p_0^B$, there is no conflict between the two, loss aversion cannot have an influence, and the Bayesian posterior is thus always optimal for all $\lambda \geq 0$. Hence, in these two cases Lemma 2 applies. However, a trade-off is present when $p_1^B < p < f^A$ or $f^A < p < p_0^B$. By Proposition 4 we know that within these two ranges the Bayesian posterior is also optimal if $p > p_1^D$ or $p < p_0^D$, so that Lemma 2 also applies if $p_1^D < p < f^A$ or $f^A < p < p_0^D$ (where both p_1^D and p_0^D depend on λ). We further know from Proposition 4 that the Bayesian posterior is not while D's prior is optimal if $p_1^B \leq p < p_1^D$ and $s = 1$ or if $p_0^D < p \leq p_0^B$ and $s = 0$. These are the two ranges we investigate in the following. In the first case where $p_1^B \leq p < p_1^D$, D demands information if and only if

$$\begin{aligned}
& \Pr(s = 1)[pu(1, 0) + (1 - p)u(0, 0)] \\
& + \Pr(s = 0)[\Pr(x = 1|p, s = 0)u(1, 0) + \Pr(x = 0|p, s = 0)u(0, 0)] - C \geq pu(1, 0) + (1 - p)u(0, 0) \\
\Leftrightarrow & [p\sigma + (1 - p)(1 - \sigma)][p\alpha + (1 - p)\beta] \\
& + [(1 - p)\sigma + p(1 - \sigma)]\left[\frac{p(1 - \sigma)}{p(1 - \sigma) + (1 - p)\sigma}\alpha + \frac{(1 - p)\sigma}{p(1 - \sigma) + (1 - p)\sigma}\beta\right] - C \geq p\alpha + (1 - p)\beta \\
\Leftrightarrow & C \leq (\alpha - \beta)(p - 1)p(2\sigma - 1),
\end{aligned}$$

where the RHS is negative (zero) if and only if $\alpha > (=) \beta$. In the second case where $p_0^D < p \leq p_0^B$, D demands information if and only if

$$\begin{aligned} & \Pr(s = 1)[\Pr(x = 1|p, s = 1)u(1, 1) + \Pr(x = 0|p, s = 1)u(0, 1)] \\ & + \Pr(s = 0)[pu(1, 1) + (1 - p)u(0, 1)] - C \geq pu(1, 1) + (1 - p)u(0, 1) \\ \Leftrightarrow & [p\sigma + (1 - p)(1 - \sigma)]\left[\frac{p\sigma}{p\sigma + (1 - p)(1 - \sigma)}\right] + [(1 - p)\sigma + p(1 - \sigma)][p] - C \geq p \\ \Leftrightarrow & C \leq p(2\sigma - 1 - p(2\sigma - 1)), \end{aligned}$$

where the RHS is positive and, if $\alpha \geq \beta$, larger than $p[\alpha(1 - \sigma) + \sigma(1 - \beta) - 1] + \beta\sigma = \bar{C}(\lambda = 0)$.

■

Proof. [Lemma 3] For any given belief $f \in Q$, Assumption 4 leads to a unique optimal action $a^*(f)$, implicitly characterized by

$$-\frac{u'(1, a^*)}{u'(0, a^*)} = \frac{1 - f}{f}. \quad (5)$$

The optimal action is unique and the first-order condition holds with equality and represents a maximum because of Conditions (i), (ii), and (iii) of Assumption 4. By the Implicit Function Theorem, we can also establish that $a^*(f)$ is strictly increasing in f .⁶⁶ By uniqueness of the optimal action, i.e., $A^*(f) = \{a^*(f)\}$, $G(A^*, f)$ is degenerate and hence $U(x, f) = u(x, a_f)$ where $a_f \equiv a^*(f)$.

We now turn to belief choice given induced actions. With y being the utility difference generated by D's prior and her potential posterior, without loss of generality we restrict attention to cases where the signal does not equal the state and use

$$y_{p,s} \equiv U(x \neq s, q(p, s)) - U(x \neq s, p) = u(x \neq s, a_{q(p,s)}) - u(x \neq s, a_p).$$

It is sufficient to only consider states unequal to the signal, because by the trade-off considered in our model we know that D's optimal belief weakly follows the signal. Since $a^*(q)$ is increasing in q and by Condition (i) of Assumption 4, we thus also know that the utility difference is weakly positive in states equal to the signal and weakly negative in states unequal to the signal. By the specification of the gain-loss function r , we thus need to consider only the latter case.

Suppose first that $s = 1$. Maximization of the target function yields a first-order condition that implicitly characterizes the unique optimal belief:

$$-\frac{U'(1, q^*(p, s = 1))}{U'(0, q^*(p, s = 1))} = \frac{\Pr(x = 0|p, s = 1)}{\Pr(x = 1|p, s = 1)} \left(1 + r'(y_{p,s=1}^*)\right), \quad (6)$$

where $y_{p,s}^* = U(x \neq s, q^*(p, s)) - U(x \neq s, p)$. Because $U'(x, f) = \frac{\partial u(x, a^*(f))}{\partial a^*(f)} \frac{\partial a^*(f)}{\partial f}$, (6) becomes

$$-\frac{u'(1, a_{q^*(p,s=1)}}{u'(0, a_{q^*(p,s=1)})} = \frac{\Pr(x = 0|p, s = 1)}{\Pr(x = 1|p, s = 1)} \left(1 + r'(y_{p,s=1}^*)\right), \quad (7)$$

⁶⁶The second-order condition is $fu''(1, a) + (1 - f)u''(0, a)$ and hence negative by Condition (ii) of Assumption 4 and the derivative of the first-order condition with respect to f is $u'(1, a) - u'(0, a)$ and positive by Condition (i).

which holds with equality and represents a maximum because of the conditions in Assumption 4.⁶⁷ The optimal belief is unique, i.e., $Q^*(p, s) = \{q^*(p, s)\}$, since $a^*(q)$ is strictly increasing in q . Plugging (5) into (7) yields

$$q^*(p, s = 1) = \frac{\Pr(x = 1|p, s = 1)}{1 + r'(y_{p,s=1}^*) \Pr(x = 0|p, s = 1)}, \quad (8)$$

which always exists by Brouwer's Fixed Point Theorem.⁶⁸ Moreover, this fixed point is unique since the RHS of (8) is monotone in q . Further, notice that $q^*(p, s = 1) \in [p, \Pr(x = 1|p, s = 1)]$: Without loss aversion, we have $r'(y) = 0 \forall y$ and hence $q^*(p, s = 1) = \Pr(x = 1|p, s = 1)$. With increasing loss aversion, $q^*(p, s = 1)$ moves towards p but never beyond (as going beyond would induce a loss in the other state of the world and is even worse in terms objective performance). Equality is reached when $q^*(p, s = 1) = p$, i.e., when

$$p = \frac{p\sigma}{p\sigma + (1-p)(1-\sigma)(1+\lambda)}.$$

Solving for λ yields $\lambda = \frac{2\sigma-1}{1-\sigma}$.

We can make an analogous argument for $s = 0$. In a nutshell, maximizing the target function yields the first-order condition

$$-\frac{u'(1, a_{q^*(p,s=0)})}{u'(0, a_{q^*(p,s=0)})} = \frac{\Pr(x = 0|p, s = 0)}{\Pr(x = 1|p, s = 0)} \times \frac{1}{1 + r'(y_{p,s=0}^*)}. \quad (9)$$

The second-order condition is again negative and plugging (5) into (9) leads us to

$$q^*(p, s = 0) = \frac{(1 + r'(y_{p,s=0}^*)) \Pr(x = 1|p, s = 0)}{1 + r'(y_{p,s=0}^*) \Pr(x = 1|p, s = 0)}, \quad (10)$$

which is again unique, i.e. $Q^*(p, s) = \{q^*(p, s)\}$, and exists by Brouwer's Fixed Point Theorem. Setting $q^*(p, s = 0) = p$ in (10) yields

$$p = \frac{p(1-\sigma)(1+\lambda)}{p(1-\sigma)(1+\lambda) + (1-p)\sigma}$$

and solving for λ yields again $\lambda = \frac{2\sigma-1}{1-\sigma}$. ■

⁶⁷The second-order condition (with q being $q(p, s = 1)$ and y being $y_{p,s=1}$) is

$$\begin{aligned} & \Pr(x = 1|p, s = 1)u''(1, a_q)a_q^2 + \Pr(x = 0|p, s = 1)a_q^2[u''(0, a_q)(1 + r'(y)) + u'(0, a_q)^2r''(y)] \\ & + a_q''[\Pr(x = 1|p, s = 1)u'(1, a_q) + \Pr(x = 0|p, s = 1)u'(0, a_q)(1 + r'(y))] \end{aligned}$$

and negative because the first line is negative by Conditions (i) and (ii) of Assumption 4 and because $r'(y) > 0$ and $r''(y) \leq 0$ for all $y < 0$. The second line equals zero by plugging in our first-order condition.

⁶⁸We have a fixed point here as y depends on q , and thus $g(q) = q$. According to Brouwer's Fixed Point Theorem this point exists since $g : [0, 1] \rightarrow [0, 1]$ is continuous.

Proof. [Lemma 4] In the following we will repeatedly use the fact that $\text{logit}(f) = \ln\left(\frac{f}{1-f}\right) = \ln(f) - \ln(1-f)$. Suppose first that $\lambda < \frac{2\sigma-1}{1-\sigma}$. We distinguish two cases. If $s = 1$, by Lemma 3 we know that

$$\begin{aligned} q^*(p, s = 1) &= \frac{p\sigma}{p\sigma + (1-p)(1-\sigma)(1+\lambda)} \quad \text{and} \\ 1 - q^*(p, s = 1) &= \frac{(1-p)(1-\sigma)(1+\lambda)}{p\sigma + (1-p)(1-\sigma)(1+\lambda)}. \end{aligned}$$

Thus,

$$\begin{aligned} \text{logit}(q^*(p, s = 1)) &= \ln\left(\frac{q^*(p, s = 1)}{1 - q^*(p, s = 1)}\right) = \ln(p\sigma) - \ln\left((1-p)(1-\sigma)(1+\lambda)\right) \\ &= \ln\left(\frac{p}{1-p}\right) + \ln\left(\frac{\sigma}{1-\sigma}\right) - \ln(1+\lambda), \end{aligned}$$

and hence

$$\text{logit}(q^*(p, s = 1)) - \text{logit}(p) = \ln\left(\frac{\sigma}{(1-\sigma)(1+\lambda)}\right).$$

In case $s = 0$, by Lemma 3 we know that

$$\begin{aligned} q^*(p, s = 0) &= \frac{p(1-\sigma)(1+\lambda)}{p(1-\sigma)(1+\lambda) + (1-p)\sigma} \quad \text{and} \\ 1 - q^*(p, s = 0) &= \frac{(1-p)\sigma}{p(1-\sigma)(1+\lambda) + (1-p)\sigma}. \end{aligned}$$

Thus,

$$\begin{aligned} \text{logit}(q^*(p, s = 0)) &= \ln\left(\frac{q^*(p, s = 0)}{1 - q^*(p, s = 0)}\right) = \ln\left(p(1-\sigma)(1+\lambda)\right) - \ln\left((1-p)\sigma\right) \\ &= \ln\left(\frac{p}{1-p}\right) + \ln\left(\frac{1-\sigma}{\sigma}\right) + \ln(1+\lambda), \end{aligned}$$

and hence

$$\text{logit}(q^*(p, s = 0)) - \text{logit}(p) = \ln\left(\frac{(1-\sigma)(1+\lambda)}{\sigma}\right).$$

Now, suppose that $\lambda \geq \frac{2\sigma-1}{1-\sigma}$. For both signals ($s = 1$ and $s = 0$), we have $q^*(p, s) = p$ (see Lemma 3). Thus,

$$\text{logit}(q^*(p, s)) = \ln\left(\frac{q^*(p, s)}{1 - q^*(p, s)}\right) = \ln\left(\frac{p}{1-p}\right) = \text{logit}(p).$$

■

Proof. [Proposition 8] We consider first the case where $\lambda < \frac{2\sigma-1}{1-\sigma}$. From Lemma 4 we know that

$$\begin{aligned} \text{logit}(q^*(p, s)) &= \text{logit}(p) + g(\sigma, \lambda), \quad \text{where} \\ g(\cdot) &= \begin{cases} \ln\left(\frac{\sigma}{(1-\sigma)(1+\lambda)}\right) & \text{if } s = 1 \equiv LLR_\lambda(s = 1|x = 1) \\ \ln\left(\frac{(1-\sigma)(1+\lambda)}{\sigma}\right) & \text{if } s = 0 \equiv LLR_\lambda(s = 0|x = 1). \end{cases} \end{aligned}$$

$LLR_\lambda(s = 1|x = 1)$ and $LLR_\lambda(s = 0|x = 1)$ refer to the “ λ -adjusted” log-likelihood ratios of the signal given $x = 1$. We can rewrite the last equation as

$$LO(x = 1|s = 1) = LO(x = 1) + LLR_\lambda(s = 1|x = 1) \quad \text{and} \quad (11)$$

$$LO(x = 1|s = 0) = LO(x = 1) + LLR_\lambda(s = 0|x = 1), \quad (12)$$

where $LO(x = 1|s = 1)$ and $LO(x = 1|s = 0)$ refer to D 's posterior log odds of $x = 1$ conditional on the observed signal, and $LO(x = 1)$ refers to the prior log odds of $x = 1$. We can again rewrite the previous expression as

$$LO(x = 1|s) = LO(x = 1) + LLR_\lambda(s|x = 1), \quad \text{where} \quad (13)$$

$$LLR_\lambda(s|x = 1) = LLR_\lambda(s = 0|x = 1) + s \times [LLR_\lambda(s = 1|x = 1) - LLR_\lambda(s = 0|x = 1)]. \quad (14)$$

The term in square brackets is the log-odds ratio of $x = 1$. To see this, note that the odds ratio of $x = 1$, associated with signal s , is

$$OR_s(x = 1) = \frac{O(x = 1|s = 1)}{O(x = 1|s = 0)},$$

where $O(x = 1|s = 1)$ and $O(x = 1|s = 0)$ refer to D 's posterior odds of $x = 1$ given the signal. Taking the log of both sides yields

$$LOR_s(x = 1) = LO(x = 1|s = 1) - LO(x = 1|s = 0) \quad (15)$$

$$= LLR_\lambda(s = 1|x = 1) - LLR_\lambda(s = 0|x = 1), \quad (16)$$

where the second equality follows from plugging (11) and (12) into (15). The RHS of (16) is equal to the term in square brackets in (14). We can therefore rewrite (14) as

$$LLR_\lambda(s|x = 1) = LLR_\lambda(s = 0|x = 1) + s \times LOR_s(x = 1). \quad (17)$$

Combining (13) and (17) yields

$$LO(x = 1|s) = LO(x = 1) + LLR_\lambda(s = 0|x = 1) + s \times LOR_s(x = 1). \quad (18)$$

And finally, combining (12) and (18) yields

$$LO(x = 1|s) = LO(x = 1|s = 0) + s \times LOR_s(x = 1),$$

so that $\text{logit}(q^*(p, s)) = \beta_0 + s \times \beta_1$.

Now, we consider the case where $\lambda \geq \frac{2\sigma-1}{1-\sigma}$. In this case,

$$\begin{aligned} LO(x = 1|s = 1) &= LO(x = 1) = LO(x = 1|s = 0) \quad \text{and thus} \\ LO(x = 1|s) &= LO(x = 1), \end{aligned}$$

so that $\text{logit}(q^*(p, s)) = \beta_0 + s \times \beta_1$ with $\beta_0 = \text{logit}(p)$ and $\beta_1 = 0$. ■

Proof. [Proposition 9] Consider any lying strategy $\tilde{\sigma}' < \sigma(\lambda)$. If the actual signal precision σ satisfies $\sigma > \sigma(\lambda)$, $\tilde{\sigma}' < \sigma(\lambda)$ performs as well as the truth-telling strategy $\tilde{\sigma} = \sigma(\lambda)$ because both strategies implement action 1. If $\sigma = \sigma(\lambda)$, $\tilde{\sigma}'$ again performs as well as $\tilde{\sigma}$ because D is now indifferent between the two actions and $\tilde{\sigma}'$ implements action 1, while $\tilde{\sigma}$ yields each action with equal probability. In contrast, if $\tilde{\sigma}' \leq \sigma < \sigma(\lambda)$, $\tilde{\sigma}$ performs better than $\tilde{\sigma}'$: the former implements action 0, which is preferred by D since $\sigma < \sigma(\lambda)$, while the latter implements either action 1 or both actions with equal probability. Finally, if $\sigma < \tilde{\sigma}'$, both strategies again perform the same because both lead to action 0. All in all, any strategy $\tilde{\sigma}' < \sigma(\lambda)$ is therefore weakly dominated by $\tilde{\sigma} = \sigma(\lambda)$. An analogous argument can be constructed for all strategies $\tilde{\sigma}'' > \sigma(\lambda)$. Again, these are weakly dominated by $\tilde{\sigma} = \sigma(\lambda)$. As a result, it is weakly dominant for D to choose $\tilde{\sigma} = \sigma(\lambda)$. ■

Proposition 10 (Conservatism) Suppose that Assumption 4 holds. For all $\lambda > 0$, updating is conservative for all priors and any signal s : If $s = 1$, $p < q^*(p, s) < \Pr(x = 1|p, s)$ for $0 < \lambda < \bar{\lambda}(\sigma)$ and $p = q^*(p, s) < \Pr(x = 1|p, s)$ for $\lambda \geq \bar{\lambda}(\sigma)$. And if $s = 0$, $p > q^*(p, s) > \Pr(x = 1|p, s)$ for $0 < \lambda < \bar{\lambda}(\sigma)$ and $p = q^*(p, s) > \Pr(x = 1|p, s)$ for $\lambda \geq \bar{\lambda}(\sigma)$

Proof. [Proposition 10] From Lemma 3 we know that in the optimum we have $g_s(q^*(p, s)) = q^*(p, s)$, where

$$\begin{aligned} g_1(q(p, s = 1)) &\equiv \frac{\Pr(x = 1|p, s = 1)}{1 + r'(y_{p,s=1}) \Pr(x = 0|p, s = 1)} = \frac{p\sigma}{p\sigma + (1-p)(1-\sigma)(1+r'(y_{p,s=1}))'} \\ g_0(q(p, s = 0)) &\equiv \frac{(1+r'(y_{p,s=0})) \Pr(x = 1|p, s = 0)}{1 + r'(y_{p,s=0}) \Pr(x = 1|p, s = 0)} = \frac{p(1-\sigma)(1+r'(y_{p,s=0}))}{p(1-\sigma)(1+r'(y_{p,s=0})) + (1-p)\sigma}. \end{aligned}$$

We first prove the case where $s = 1$. Notice that we do not need to consider the case $q(p, s = 1) < p$ as it is worse than $q(p, s = 1) = p$ in terms of objective performance as well as loss aversion (because $a^*(f)$ is strictly increasing in f). Suppose that $q(p, s = 1) = p$. Then, $r'(y_{p,s=1}) = \lambda$ and hence $g_1(q(p, s = 1) = p) < (=) \Pr(x = 1|p, s = 1)$ for $\lambda > (=) 0$. Also, $g_1(q(p, s = 1) = p) > (=) p$ for $\lambda < (=) \frac{2\sigma-1}{1-\sigma}$. For the cases where $q(p, s = 1) > p$ we can look at

$$\begin{aligned} g_1'(q(p, s = 1)) &= -\frac{p\sigma(1-p)(1-\sigma)r''(y_{p,s=1})u'(0, a_{q(p,s=1)})a'_{q(p,s=1)}}{[p\sigma + (1-p)(1-\sigma)(1+r'(y_{p,s=1}))]^2} \\ &< (=) 0 \quad \text{for} \quad r''(y_{p,s=1}) < (=) 0. \end{aligned}$$

Notice first that if $g_1(q(p, s = 1) = p) > p$, we have $g_1(q^*(p, s = 1)) > p$ because $g_1'(q(p, s = 1) = p) > -\infty$. On the other hand, if $g_1(q(p, s = 1) = p) = p$, we have $g_1(q^*(p, s = 1)) = p$ because

$g'_1(q(p, s = 1)) \leq 0$. Thus, we have $q^*(p, s = 1) > (=) p$ for $\lambda < (\geq) \frac{2\sigma-1}{1-\sigma}$. Because signals are informative, i.e., $\sigma > 1/2$, we always have $\Pr(x = 1|p, s = 1) > p$. Secondly, we can distinguish the following two cases. First, if $\lambda = 0$ (and thus $r'(y) = 0 \forall y$), we have $q^*(p, s = 1) = \Pr(x = 1|p, s = 1)$ because $g_1(q(p, s = 1) = p) = \Pr(x = 1|p, s = 1)$ and $g'_1(q(p, s = 1)) = 0 \forall q(p, s = 1)$. Second, if $\lambda > 0$, we have $q^*(p, s = 1) < \Pr(x = 1|p, s = 1)$ because $g_1(q(p, s = 1) = p) < \Pr(x = 1|p, s = 1)$ and $g'_1(q(p, s = 1)) \leq 0$. Figure 4 in the proof of Proposition 11 illustrates these arguments.

We next prove the case where $s = 0$. We do not need to consider the case $q(p, s = 0) > p$ as it is worse than $q(p, s = 0) = p$ in terms of objective performance as well as loss aversion ($a^*(f)$ is strictly increasing in f). Suppose that $q(p, s = 0) = p$. Then, $r'(y_{p,s=0}) = \lambda$ and hence $g_0(q(p, s = 0) = p) > (=) \Pr(x = 1|p, s = 0)$ for $\lambda > (=) 0$. Also, $g_0(q(p, s = 0) = p) < (=) p$ for $\lambda < (\geq) \frac{2\sigma-1}{1-\sigma}$. For the cases where $q(p, s = 0) < p$ we can look at

$$\begin{aligned} g'_0(q(p, s = 0)) &= \frac{\sigma p(1-p)(1-\sigma)r''(y_{p,s=0})u'(1, a_{q(p,s=0)})a'_{q(p,s=0)}}{[p(1-\sigma)(1+r'(y_{p,s=0})) + (1-p)\sigma]^2} \\ &< (=) 0 \quad \text{for } r''(y_{p,s=0}) < (=) 0. \end{aligned}$$

First, notice that if $g_0(q(p, s = 0) = p) < p$, we have $g_0(q^*(p, s = 0)) < p$ because $g'_0(q(p, s = 0) = p) > -\infty$. On the other hand, if $g_0(q(p, s = 0) = p) = p$, we have $g_0(q^*(p, s = 0)) = p$ because $g'_0(q(p, s = 0)) \leq 0$. Thus, we have $q^*(p, s = 0) < (=) p$ for $\lambda < (\geq) \frac{2\sigma-1}{1-\sigma}$. Because signals are informative, i.e., $\sigma > 1/2$, we always have $\Pr(x = 1|p, s = 0) < p$. Second, we distinguish the following two cases. First, if $\lambda = 0$ (and thus $r'(y) = 0 \forall y$), we have $q^*(p, s = 0) = \Pr(x = 1|p, s = 0)$ because $g_0(q(p, s = 0) = p) = \Pr(x = 1|p, s = 0)$ and $g'_0(q(p, s = 0)) = 0 \forall q(p, s = 0)$. Second, if $\lambda > 0$, we have $q^*(p, s = 0) > \Pr(x = 1|p, s = 0)$ because $g_0(q(p, s = 0) = p) < \Pr(x = 1|p, s = 0)$ and $g'_0(q(p, s = 0)) \leq 0$. Again, Figure 4 in the proof of Proposition 11 illustrates. ■

Corollary 4 (Maximal Conservatism) *Suppose that Assumptions 1 and 4 hold and $\lambda > 0$. If $s = 1$, conservatism is maximal for some $p < 1/2$, i.e., $\arg \max_{p \in (0,1)} \Pr(x = 1|p, s = 1) - q^* < 1/2$. And if $s = 0$, conservatism is maximal for some $p > 1/2$, i.e., $\arg \max_{p \in (0,1)} q^* - \Pr(x = 1|p, s = 1) > 1/2$.*

For the remaining results, it is sometimes useful to rely on two simplifying assumptions to keep the analysis tractable. The first assumption abstracts away from third-order effects:

Assumption 5 (Linear Marginal (Loss) Utility) $u'''(\cdot, a) = 0 \forall a$ and $r'''(y) = 0 \forall y < 0$.

The second assumption restricts attention to cases where the utility loss of taking a worse action is weakly larger in the good than bad state:

Assumption 6 (Weakly More Concavity in Good State) $u''(0, a) \geq u''(1, a)$.

From Lemma 3 and Proposition 10 it follows that updating is symmetric for $\lambda \geq \frac{2\sigma-1}{1-\sigma}$ or $\lambda = 0$, because the prior is chosen in the former and the Bayesian posterior in the latter case. Proposition 11 therefore restricts attention to the remaining loss aversion values.

Proposition 11 (Asymmetry) *Suppose that Assumptions 4 to 6 hold and $0 < \lambda < \bar{\lambda}(\sigma)$. If $r''(y) = 0 \forall y < 0$, updating is symmetric for all priors, i.e., good news are updated from p as conservatively as bad*

news from $1 - p$ for all p . If $r''(y) < 0 \forall y < 0$, updating can be asymmetric: good news are updated from p less conservatively than bad news from $1 - p$ for all p if $u''(1, a) < u''(0, a)$ and provided that σ is sufficiently large.

Proof. [Proposition 11] Updating is symmetric if and only if

$$q^*(p, s = 1) - p = 1 - p - q^*(1 - p, s = 0) \Leftrightarrow q^*(p, s = 1) = 1 - q^*(1 - p, s = 0).$$

By the proof of Lemma 3 we know that

$$q^*(p, s = 0) = \frac{p(1 - \sigma) \left(1 + r'(y_{p,s=0}^*)\right)}{p(1 - \sigma) \left(1 + r'(y_{p,s=0}^*)\right) + (1 - p)\sigma}.$$

We can therefore establish that

$$1 - q^*(1 - p, s = 0) = \frac{p\sigma}{p\sigma + (1 - p)(1 - \sigma) \left(1 + r'(y_{1-p,s=0}^*)\right)} = q^*(p, s = 1),$$

where the last equality again follows from the proof of Lemma 3 if $r'(y_{1-p,s=0}^*) = r'(y_{p,s=1}^*)$, which is always fulfilled when $r''(y) = 0 \forall y < 0$ because then $r'(y) = \lambda \forall y < 0$.

We next turn to asymmetric updating, by which we mean that good news are updated less conservatively than bad news. More formally, we say that updating is asymmetric if and only if

$$q^*(p, s = 1) - p > 1 - p - q^*(1 - p, s = 0) \Leftrightarrow q^*(p, s = 1) > 1 - q^*(1 - p, s = 0).$$

Because

$$\begin{aligned} q^*(p, s = 1) &= \frac{p\sigma}{p\sigma + (1 - p)(1 - \sigma) \left(1 + r'(y_{p,s=1}^*)\right)} \quad \text{and} \\ 1 - q^*(1 - p, s = 0) &= \frac{p\sigma}{p\sigma + (1 - p)(1 - \sigma) \left(1 + r'(y_{1-p,s=0}^*)\right)}, \end{aligned}$$

by the concavity of $r(\cdot)$ and because $y_{p,s=1}$ and $y_{1-p,s=0}$ are weakly negative, we have that

$$q^*(p, s = 1) > 1 - q^*(1 - p, s = 0) \Leftrightarrow y_{p,s=1}^* > y_{1-p,s=0}^*.$$

Hence, in order to get asymmetric updating we require the utility difference from optimally taking an action under the chosen posterior (and further away from the state than it was the case under the prior) and under the prior to be larger in the good than in the bad state.


In order to see when this is the case, we need to compare the two fixed points from above:

$$q^*(p, s = 1) = \frac{p\sigma}{p\sigma + (1 - p)(1 - \sigma) \left(1 + r'(y_{p,s=1}^*)\right)} \equiv \text{FP1}, \quad (19)$$

$$q^*(1 - p, s = 0) = 1 - \frac{p\sigma}{p\sigma + (1 - p)(1 - \sigma) \left(1 + r'(y_{1-p,s=0}^*)\right)} \equiv \text{FP0}. \quad (20)$$

In the following, we denote the RHS of equality (19) by $g_1(q(p, s = 1))$ and of equality (20) by $g_0(q(1 - p, s = 0))$. Good news are updated less conservatively than bad news if the fixed point FP1 is further away from p than the fixed point FP0 from $1 - p$, because then we have $q^*(p, s = 1) > 1 - q^*(1 - p, s = 0)$. Figure 4 exemplifies.

Figure 4: Comparing Fixed Points



Notes: The right curve represents $g_1(q(p, s = 1))$ and the left curve $g_0(q(1 - p, s = 0))$ within the range of interest. The dots at the end of the curves represent the respective fixed points.

By Proposition 10 we know that the optimal solution satisfies $p < q^*(p, s = 1) < \Pr(x = 1|p, s = 1)$ and $\Pr(x = 1|1 - p, s = 0) < q^*(1 - p, s = 0) < p$. Notice that $y_{p,s=1}|_{q=p} = y_{1-p,s=0}|_{q=1-p} = 0$. By the proof of Proposition 10 and since $0 < \lambda < \frac{2\sigma-1}{1-\sigma}$, we also know that $\Pr(x = 1|p, s = 1) > g_1(q(p, s = 1) = p) > p$ and $\Pr(x = 1|1 - p, s = 0) < g_0(q(1 - p, s = 0) = 1 - p) < 1 - p$. Moreover, since $y_{p,s=1}|_{q=p} = y_{1-p,s=0}|_{q=1-p}$, we have that $g_0(q(1 - p, s = 0) = 1 - p) = 1 - g_1(q(p, s = 1) = p)$. Therefore, a sufficient condition for having FP1 further away from p than FP0 from $1 - p$ for all p is that the negative slope of $g_0(q(1 - p, s = 0) < 1 - p)$ is absolutely larger than the negative slope of $g_1(q(p, s = 1) > p)$ for all $q(p, s = 1)$ and $q(1 - p, s = 0)$ that satisfy $q(p, s = 1) - p = 1 - p - q(1 - p, s = 0)$ until $g_0(\cdot)$ has reached the identity line (which happens faster than for $g_1(\cdot)$ under asymmetry).

The slopes of $g_1(q(p, s = 1))$ and $g_0(q(1 - p, s = 0))$ with respect to q are

$$g_1'(q(p, s = 1)) = -\frac{p\sigma(1-p)(1-\sigma)r''(y_{p,s=1})u'(0, a_{q(p,s=1)})a'_{q(p,s=1)}}{[p\sigma + (1-p)(1-\sigma)(1+r'(y_{p,s=1}))]^2} \quad \text{and}$$

$$g_0'(q(1-p, s = 0)) = \frac{p\sigma(1-p)(1-\sigma)r''(y_{1-p,s=0})u'(1, a_{q(1-p,s=0)})a'_{q(1-p,s=0)}}{[p\sigma + (1-p)(1-\sigma)(1+r'(y_{1-p,s=0}))]^2},$$

which are both weakly negative. More specifically, both slopes are zero for $r''(y) = 0 \forall y < 0$ and strictly negative when $r''(y) < 0 \forall y < 0$. Because equal slopes imply symmetric updating, slopes of zero also imply symmetric updating. In the following, we therefore focus on the case where $r''(y) < 0 \forall y < 0$.

Notice first that $g'_1(q(p, s = 1)) > (=) g'_0(q(1 - p, s = 0))$ if and only if

$$\frac{-u'(0, a_{q(p,s=1)})a'_{q(p,s=1)}}{u'(1, a_{q(1-p,s=0)})a'_{q(1-p,s=0)}} < (=) \frac{r''(y_{1-p,s=0}) [p\sigma + (1-p)(1-\sigma) (1 + r'(y_{p,s=1}))]^2}{r''(y_{p,s=1}) [p\sigma + (1-p)(1-\sigma) (1 + r'(y_{1-p,s=0}))]^2}, \quad (21)$$

where the RHS of (21) is equal to 1 at $q(p, s = 1) = p$ and $q(1 - p, s = 0) = 1 - p$ since $y_{p,s=1}|_{q=p} = y_{1-p,s=0}|_{q=1-p} = 0$. Moreover, under the assumption that $r'''(y) = 0 \forall y < 0$ (Assumption 5), we have that $r''(y_{1-p,s=0}) = r''(y_{p,s=1})$. So, in the limit when σ converges to 1, the RHS of (21) is equal to 1, i.e., $\lim_{\sigma \rightarrow 1} RHS(21) = 1$. Hence, when σ is sufficiently large,

$$-u'(0, a_{q(p,s=1)})a'_{q(p,s=1)} < u'(1, a_{q(1-p,s=0)})a'_{q(1-p,s=0)} \quad (22)$$

implies $g'_1(q(p, s = 1)) > g'_0(q(1 - p, s = 0))$. We can use the Implicit Function Theorem on (5) in the proof of Lemma 3 to derive

$$a'_f = -\frac{u'(1, a_f) - u'(0, a_f)}{f u''(1, a_f) + (1 - f) u''(0, a_f)} > 0.$$

Plugging it into (22) and evaluating at $q(p, s = 1) = p$ and $q(1 - p, s = 0) = 1 - p$ yields

$$\begin{aligned} & u'(1, a_{1-p}) [u'(1, a_{1-p}) - u'(0, a_{1-p})] [p u''(1, a_p) + (1 - p) u''(0, a_p)] \\ < & -u'(0, a_p) [u'(1, a_p) - u'(0, a_p)] [(1 - p) u''(1, a_{1-p}) + p u''(0, a_{1-p})], \end{aligned} \quad (23)$$

with both sides of (23) being negative. Imposing $u'''(x, a) = 0 \forall a$ (Assumption 5) and using (5) in the proof of Lemma 3, we can establish that the left-hand side (LHS) of (23) is smaller (equal) than its RHS for $p > (=) 1/2$ if $u''(1, a) < u''(0, a)$. When $u''(1, a) = u''(0, a)$, both sides of (23) are equal. Thus, for $u''(1, a) < u''(0, a)$, we need to distinguish three cases: $p > 1/2$, $p = 1/2$, and $p < 1/2$.

Suppose first that $p > 1/2$: We know that $0 > g'_1(q(p, s = 1) = p) > g'_0(q(1 - p|s = 0) = 1 - p)$. What happens as $q(p|s = 1)$ moves away from p and $q(1 - p|s = 0)$ from $1 - p$ by the same amount? Since the difference between the LHS and the RHS of (23) increases in p for $p > 1/2$, we also know that the difference

$$u'(1, a_{q(1-p,s=0)})a'_{q(1-p,s=0)} - \left(-u'(0, a_{q(p,s=1)})a'_{q(p,s=1)} \right)$$

is positive and increases as $q(\cdot)$ symmetrically moves away from p and $1 - p$. Hence, this effect pushes the observed difference at $q(p, s = 1) = p$ and $q(1 - p, s = 0) = 1 - p$ between $g'_1(q(p, s = 1))$ and $g'_0(q(1 - p, s = 0))$ even further apart as we move away. Next, suppose that $p = 1/2$: In this case we know that $0 > g'_1(q(p, s = 1) = p) = g'_0(q(1 - p, s = 0) = 1 - p)$. But, by the same argument as before, symmetrically moving away from p and $1 - p$ implies $g'_1(q(p, s = 1)) > g'_0(q(1 - p, s = 0))$. Last, suppose that $p < 1/2$: In this case, we know that $g'_1(q(p|s = 1) = p) < g'_0(q(1 - p, s = 0) = 1 - p) < 0$. What happens as $q(p, s = 1)$ moves away from p and $q(1 - p, s = 0)$ from $1 - p$? The difference between the LHS and the RHS of (23) becomes smaller and when crossing $1/2$ even changes its sign. Hence, $g'_1(q(p, s = 1))$ will be closer to and, eventually, be larger than $g'_0(q(1 - p, s = 0))$. Since these changes are symmetric around

$1/2$, a sufficiently large σ can always achieve that, across all $q(\cdot)$ in the relevant range, the average $g'_1(q(p, s = 1))$ is larger than the average $g'_0(q(1 - p, s = 0))$. ■

Proposition 12 (Information Preferences) *We distinguish the following three cases:*

- (i) *Suppose that Assumption 4 holds but $\sigma = 1$. For all $\lambda \geq 0$, the decision maker is strictly information seeking for all priors, i.e., $\bar{C} > 0$;*
- (ii) *Suppose that Assumptions 4 to 6 hold. For $\lambda = 0$, the decision maker is strictly information seeking for all priors, i.e., $\bar{C} > 0$;*
- (iii) *Suppose that Assumption 4 holds and $\lambda > 0$. If $0 < \lambda < \bar{\lambda}(\sigma)$, the decision maker can be strictly information averse for low priors, i.e., $\bar{C} < (=) 0$ if and only if $u(1, 0) > (=) u(0, 0)$, but is strictly information seeking for large priors, i.e., $\bar{C} > 0$. And if $\lambda \geq \bar{\lambda}(\sigma)$, the decision maker is information neutral for all priors, i.e., $\bar{C} = 0$.*

Proof. [Proposition 12] To simplify notation we use $a_{q^*(p, s=1)} \equiv a_1$ and $a_{q^*(p, s=0)} \equiv a_0$, where $a_f \equiv a^*(f)$. We prove each case of Proposition 12 separately and start with case (i). Since $\sigma = 1$, by the proof of Lemma 3 we have that $q^*(p, s = 1) = 1$ and $a_1 = 1$ as well as $q^*(p, s = 0) = 0$ and $a_0 = 0$. Hence, D demands information if and only if

$$\begin{aligned} & \Pr(s = 1)u(1, 1) + \Pr(s = 0)u(0, 0) - C \geq pu(1, a_p) + (1 - p)u(0, a_p) \\ \Leftrightarrow & C \leq p[u(1, 1) - u(1, a_p)] + (1 - p)[u(0, 0) - u(0, a_p)], \end{aligned}$$

where the RHS is positive for $p \in (0, 1)$ as then $a_p \in (0, 1)$ (see the proof of Lemma 3).

We next prove case (ii). From Lemma 3 we know that $q^*(p, s = 1) = \Pr(x = 1|p, s = 1)$ and $q^*(p, s = 0) = \Pr(x = 1|p, s = 0)$. Hence, D demands information if and only if

$$\begin{aligned} & \Pr(s = 1) \left[\Pr(x = 1|p, s = 1)u(1, a_1) + \Pr(x = 0|p, s = 1)u(0, a_1) \right] \\ & + \Pr(s = 0) \left[\Pr(x = 1|p, s = 0)u(1, a_0) + \Pr(x = 0|p, s = 0)u(0, a_0) \right] - C \\ & \geq pu(1, a_p) + (1 - p)u(0, a_p) \\ \Leftrightarrow & C \leq p \left[\sigma u(1, a_1) + (1 - \sigma)u(1, a_0) - u(1, a_p) \right] + (1 - p) \left[\sigma u(0, a_0) + (1 - \sigma)u(0, a_1) - u(0, a_p) \right] \end{aligned}$$

and we need to show that the RHS (denoted \bar{C}) is positive for all $p \in (0, 1)$. Notice first that

$$\lim_{\sigma \rightarrow 1} \bar{C} = p \left[u(1, 1) - u(1, a_p) \right] + (1 - p) \left[u(0, 0) - u(0, a_p) \right],$$

where the RHS is positive for all p by the proof of case (i) above. Notice also that

$$\begin{aligned} \lim_{\sigma \rightarrow 0.5} \bar{C} &= p \left[\frac{1}{2}u(1, a_1) + \frac{1}{2}u(1, a_0) - u(1, a_p) \right] + (1 - p) \left[\frac{1}{2}u(0, a_0) + \frac{1}{2}u(0, a_1) - u(0, a_p) \right] \\ &= 0 \forall p, \end{aligned}$$

where the second equality follows from the fact that in the limit we have $q^*(p, s = 1) = p = q^*(p, s = 0)$ and thus $a_1 = a_p = a_0$. The last step is to show that $\frac{d\bar{C}}{d\sigma} > 0 \forall p$. Notice that

$$\begin{aligned} \frac{d\bar{C}}{d\sigma} &= p \left[u(1, a_1) - u(1, a_0) \right] + (1 - p) \left[u(0, a_0) - u(0, a_1) \right] \\ &\quad + p^2(1 - p) \left[\frac{\sigma}{(1 - \sigma + (2\sigma - 1)p)^2} a'_1 u'(1, a_1) - \frac{1 - \sigma}{(p + \sigma - 2p\sigma)^2} a'_0 u'(1, a_0) \right] \\ &\quad + p(1 - p)^2 \left[\frac{1 - \sigma}{(1 - \sigma + (2\sigma - 1)p)^2} a'_1 u'(0, a_1) - \frac{\sigma}{(p + \sigma - 2p\sigma)^2} a'_0 u'(0, a_0) \right]. \end{aligned}$$

It is easy to see that the first two terms are positive for all p , but it is less obvious for the last two terms. Consider the third term first and suppose in negation that

$$p^2(1 - p) \frac{1 - \sigma}{(p + \sigma - 2p\sigma)^2} a'_0 u'(1, a_0) > p^2(1 - p) \frac{\sigma}{(1 - \sigma + (2\sigma - 1)p)^2} a'_1 u'(1, a_1) \quad (24)$$

We can use the Implicit Function Theorem on (5) in the proof of Lemma 3 to derive

$$a'_f = - \frac{u'(1, a_f) - u'(0, a_f)}{f u''(1, a_f) + (1 - f) u''(0, a_f)} > 0.$$

Plugging it into (24) and imposing $u'''(\cdot, a) = 0 \forall a$ (Assumption 5) yields

$$\begin{aligned} &\frac{1 - \sigma}{(p + \sigma - 2p\sigma)^2} u'(1, a_0) \left[u'(1, a_0) - u'(0, a_0) \right] \\ &> \frac{\sigma}{(1 - \sigma + (2\sigma - 1)p)^2} u'(1, a_1) \left[u'(1, a_1) - u'(0, a_1) \right]. \end{aligned} \quad (25)$$

When we restrict attention to $u''(0, a) \geq u''(1, a)$, it can be shown that $\frac{dLHS(25)}{d\sigma} < (=, >) 0$ for $p < (=, >) 1/2$ whereas $\frac{dRHS(25)}{d\sigma} > (=, <) 0$ for $p < (=, >) 1/2$. Suppose first that $p < 1/2$. Then, the difference in (25) should be largest for $\sigma \rightarrow 0.5$, but since then $q^*(p, s = 1) = p = q^*(p, s = 0)$ we get $a_1 = a_p = a_0$ and thus

$$\lim_{\sigma \rightarrow 0.5} LHS(25) = 2 > 2 = \lim_{\sigma \rightarrow 0.5} RHS(25),$$

a contradiction. Next, suppose that $p = 1/2$. Then, the difference in (25) remains constant for all σ and it is sufficient to look at the case where $\sigma \rightarrow 0.5$. Thus, we get the same contradiction as before. Last, suppose that $p > 1/2$. In this case, the difference in (25) should be largest for $\sigma \rightarrow 1$, but since then $q^*(p, s = 1) = 1$ and $q^*(p, s = 0) = 0$, we get that $a_1 = 1$ and $a_0 = 0$ and thus

$$\lim_{\sigma \rightarrow 1} LHS(25) = 0 > 0 = \lim_{\sigma \rightarrow 1} RHS(25),$$

a contradiction. It follows that the third term of $\frac{d\bar{C}}{d\sigma}$ is weakly positive for all p . The argument for the last term of $\frac{d\bar{C}}{d\sigma}$ is similar and establishes that also the fourth term is weakly positive for all p . Since the first two terms are strictly positive for all p , it follows that $\frac{d\bar{C}}{d\sigma} > 0 \forall p$.

Last, we prove case (iii). We start with $0 < \lambda < \frac{2\sigma-1}{1-\sigma}$. D demands information if and only if

$$\begin{aligned}
& \Pr(s=1) \left[q^*(p, s=1)u(1, a_1) + (1 - q^*(p, s=1))u(0, a_1) \right] \\
& + \Pr(s=0) \left[q^*(p, s=0)u(1, a_0) + (1 - q^*(p, s=0))u(0, a_0) \right] - C \\
\geq & pu(1, a_p) + (1-p)u(0, a_p) \\
\Leftrightarrow & C \leq \Pr(s=1) \left[q^*(p, s=1)u(1, a_1) + (1 - q^*(p, s=1))u(0, a_1) - \frac{p}{\Pr(s=1)}u(1, a_p) \right] \\
& + \Pr(s=0) \left[q^*(p, s=0)u(1, a_0) + (1 - q^*(p, s=0))u(0, a_0) - \frac{1-p}{\Pr(s=0)}u(0, a_p) \right],
\end{aligned}$$

where the RHS is denoted \bar{C} . We can use Lemma 3 to determine the exact value of \bar{C} . After plugging in, we notice first that

$$\begin{aligned}
\lim_{p \rightarrow 0} \bar{C} &= (1-\sigma) \left[0u(1,0) + 1u(0,0) - 0 \right] + \sigma \left[0u(1,0) + 1u(0,0) - \frac{1}{\sigma}u(0,0) \right] \\
&= u(0,0)(1-\sigma + \sigma - 1) = 0
\end{aligned}$$

and also that

$$\begin{aligned}
\lim_{p \rightarrow 1} \bar{C} &= \sigma \left[1u(1,1) + 0u(0,1) - \frac{1}{\sigma}u(1,1) \right] + (1-\sigma) \left[1u(1,1) + 0u(0,1) - 0 \right] \\
&= u(1,1)(\sigma - 1 + 1 - \sigma) = 0.
\end{aligned}$$

Secondly, we have

$$\begin{aligned}
\left. \frac{d\bar{C}}{dp} \right|_{p=0} &= \frac{1}{1+\lambda} \left[\lambda \left(1 - 2\sigma + \lambda(1-\sigma) \right) \left(u(1,0) - u(0,0) \right) \right] \\
&< (=) 0 \Leftrightarrow u(1,0) > (=) u(0,0),
\end{aligned}$$

where the second line follows because $\lambda > 0$ and because $1 - 2\sigma + \lambda(1-\sigma) < 0 \Leftrightarrow \lambda < \frac{2\sigma-1}{1-\sigma}$. Hence, $\bar{C} < (=) 0$ for small p if and only if $u(1,0) > (=) u(0,0)$. Further, we have

$$\begin{aligned}
\left. \frac{d\bar{C}}{dp} \right|_{p=1} &= \frac{1}{1+\lambda} \left[\lambda \left(1 - 2\sigma + \lambda(1-\sigma) \right) \left(u(1,1) - u(0,1) \right) \right] \\
&< (=) 0 \Leftrightarrow u(1,1) > (=) u(0,1),
\end{aligned}$$

where the second line follows by the same argument as before. By Assumption 4 we know that $u(1,1) > u(0,1)$ and we therefore have that $\bar{C} > 0$ for large p .

Now, assume that $\lambda \geq \frac{2\sigma-1}{1-\sigma}$. Then, D demands information if and only if

$$\begin{aligned}
& \Pr(s=1) \left[pu(1, a_p) + (1-p)u(0, a_p) \right] + \Pr(s=0) \left[pu(1, a_p) + (1-p)u(0, a_p) \right] - C \\
\geq & pu(1, a_p) + (1-p)u(0, a_p) \\
\Leftrightarrow & C \leq 0,
\end{aligned}$$

where the RHS is denoted \bar{C} . ■

References

- Abarbanell, Jeffrey S. and Victor L. Bernard**, "Tests of Analysts' Overreaction/Underreaction to Earnings Information as an Explanation for Anomalous Stock Price Behavior," *Journal of Finance*, 1992, 47 (3), 1181–1207.
- Abelson, Robert P.**, "Beliefs Are Like Possessions," *Journal for the Theory of Social Behaviour*, 10 1986, 16 (3), 223–250.
- Akerlof, George A. and William T. Dickens**, "The Economic Consequences of Cognitive Dissonance," *American Economic Review*, 1982, 72 (3), 307–319.
- Ambuehl, Sandro and Shengwu Li**, "Belief Updating and the Demand for Information," 2015. Working Paper.
- Ball, Ray and Philip Brown**, "An Empirical Evaluation of Accounting Income Numbers," *Journal of Accounting Research*, 1968, pp. 159–178.
- Barberis, Nicholas, Andrei Shleifer, and Robert Vishny**, "A model of Investor Sentiment," *Journal of Financial Economics*, 1998, 49 (3), 307–343.
- Barigozzi, Francesca and Rosella Levaggi**, "Emotional Decision-Makers and Anomalous Attitudes towards Information," *Journal of Risk and Uncertainty*, 2010, 40 (3), 255–280.
- Bell, David E.**, "Regret in Decision Making under Uncertainty," *Operations Research*, 1982, 30 (5), 961–981.
- Bénabou, Roland**, "Groupthink: Collective Delusions in Organizations and Markets," *Review of Economic Studies*, 2013, 80, 429–462.
- **and Jean Tirole**, "Self-Confidence and Personal Motivation," *Quarterly Journal of Economics*, 2002, 117 (3), 871–915.
- **and —**, "Willpower and Personal Rules," *Journal of Political Economy*, 2004, 112 (4), 848–887.
- **and —**, "Belief in a Just World and Redistributive Politics," *Quarterly Journal of Economics*, 2006, 121 (2), 699–746.
- **and —**, "Identity, Morals, and Taboos: Beliefs as Assets," *Quarterly Journal of Economics*, 2011, 126, 805–855.
- Benjamin, Daniel J., Matthew Rabin, and Collin Raymond**, "A Model of Non-Belief in the Law of Large Numbers," *Journal of the European Economic Association*, 2016, 14 (2), 515–544.
- Benoît, Jean-Pierre and Juan Dubra**, "Apparent Overconfidence," *Econometrica*, 2011, 79 (5), 1591–1625.
- Bernard, Victor L.**, "Stock Price Reactions to Earnings Announcements: A Summary of Recent Anomalous Evidence and Possible Explanations," in Richard H. Thaler, ed., *Advances in Behavioral Finance*, Russell Sage Foundation, 1993, pp. 303–340.
- **and Jacob K. Thomas**, "Post-Earnings-Announcement Drift: Delayed Price Response or Risk Premium?," *Journal of Accounting Research*, 1989, pp. 1–36.
- **and —**, "Evidence That Stock Prices Do Not Fully Reflect the Implications of Current Earnings for Future Earnings," *Journal of Accounting and Economics*, 1990, 13 (4), 305–340.
- Bodoh-Creed, Aaron L.**, "The Economics of Base-Rate Neglect," 2009. Working Paper.
- Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer**, "Stereotypes," *Quarterly Journal of Economics*, Forthcoming.
- **, Nicola Gennaioli, and Andrei Shleifer**, "Diagnostic Expectations and Credit Cycles," 2016. Working Paper.
- Bouchaud, Jean-Philippe, Philipp Krueger, Augustin Landier, and David Thesmar**, "Sticky Expectations and Stock Market Anomalies," 2016. Working Paper.
- Bowman, David, Deborah Minehart, and Matthew Rabin**, "Loss Aversion in a Consumption-Savings Model," *Journal of Economic Behavior and Organization*, 1999, 38 (2), 155–178.
- Brunnermeier, Markus K. and Jonathan A. Parker**, "Optimal Expectations," *American Economic Review*, 2005, 95 (4), 1092–1118.
- **, Christian Gollier, and Jonathan A. Parker**, "Optimal Beliefs, Asset Prices, and the Preference for Skewed Returns," *American Economic Review*, 2007, 97 (2), 159–165.

Burks, Stephen V., Jeffrey P. Carpenter, Lorenz Goette, and Aldo Rustichini, "Overconfidence and Social Signalling," *Review of Economic Studies*, 2013, 80 (3), 949–983.

Buser, Thomas, Leonie Gerhards, and Joël J. van der Weele, "Measuring Responsiveness to Feedback as a Personal Trait," 2016. Working Paper.

Caplin, Andrew and John Leahy, "Psychological Expected Utility Theory and Anticipatory Feelings," *Quarterly Journal of Economics*, 2001, 116 (1), 55–79.

— **and —**, "The Supply of Information by a Concerned Expert," *Economic Journal*, 2004, 114 (497), 487–505.

Carrillo, Juan D. and Thomas Mariotti, "Strategic Ignorance as a Self-Disciplining Device," *Review of Economic Studies*, 2000, 67 (3), 529–544.

Chan, Wesley S., "Stock Price Reaction to News and No-News: Drift and Reversal After Headlines," *Journal of Financial Economics*, 2003, 70 (2), 223–260.

Choi, Darwin and Dong Lou, "A Test of the Self-Serving Attribution Bias: Evidence from Mutual Funds," 2010. Working Paper.

— **and Sam K. Hui,** "The Role of Surprise: Understanding Overreaction and Underreaction to Unanticipated Events Using In-Play Soccer Betting Market," *Journal of Economic Behavior & Organization*, 2014, 107, 614–629.

— **, Bige Kahraman, and Abhiroop Mukherjee,** "Learning About Mutual Fund Managers," *Journal of Finance*, Forthcoming.

Cialdini, Robert B., *Influence: The Psychology of Persuasion*, New York, USA: Harper Collins New York, 1984.

Coutts, Alexander, "Good News and Bad News are Still News: Experimental Evidence on Belief Updating," 2016. Working Paper.

DeBondt, Werner F.M. and Richard H. Thaler, "Do Security Analysts Overreact?," *American Economic Review*, 1990, 80 (2), 52–57.

— **and Richard Thaler,** "Does the Stock Market Overreact?," *Journal of Finance*, 1985, 40 (3), 793–805.

Edwards, Ward, "Conservatism in Human Information Processing," in Benjamin Kleinmuntz, ed., *Formal Representation of Human Judgment*, Wiley New York, 1968, pp. 17–52.

Eil, David and Justin M. Rao, "The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself," *American Economic Journal: Microeconomics*, 2011, 3 (2), 114–38.

Eliaz, Kfir and Ran Spiegler, "Can Anticipatory Feelings Explain Anomalous Choices of Information Sources?," *Games and Economic Behavior*, 2006, 56 (1), 87–104.

Epley, Nicholas and Thomas Gilovich, "Putting Adjustment Back in the Anchoring and Adjustment Heuristic: Differential Processing of Self-Generated and Experimenter-Provided Anchors," *Psychological Science*, 2001, 12 (5), 391–396.

— **and —**, "The Anchoring-And-Adjustment Heuristic – Why the Adjustments are Insufficient," *Psychological Science*, 2006, 17 (4), 311–318.

Epstein, Larry G., Jawwad Noor, and Alvaro Sandroni, "Non-Bayesian Learning," *BE Journal of Theoretical Economics*, 2010, 10 (1).

Ertac, Seda, "Does Self-Relevance Affect Information Processing? Experimental Evidence on the Response to Performance and Non-Performance Feedback," *Journal of Economic Behavior and Organization*, 2011, 80 (3), 532–545.

Eyster, Erik, "Rationalizing the Past: A Taste for Consistency," 2002. Working Paper.

Falk, Armin and Florian Zimmermann, "Information Processing and Commitment," 2015. Working Paper.

Friedman, Daniel, "Monty Hall's Three Doors: Construction and Deconstruction of a Choice Anomaly," *American Economic Review*, 1998, 88 (4), 933–946.

Fryer, Jr. Roland G., Philipp Harms, and Matthew O. Jackson, "Updating Beliefs when Evidence is Open to Interpretation: Implications for Bias and Polarization," 2015. Working Paper.

- Gennaioli, Nicola and Andrei Shleifer**, "What Comes to Mind," *Quarterly Journal of Economics*, 2010, 125 (4), 1399–1433.
- , —, and **Robert Vishny**, "Neglected Risks: The Psychology of Financial Crises," *American Economic Review: Papers and Proceedings*, 2015, 105 (5), 310–314.
- Gilovich, Thomas and Victoria H. Medvec**, "The Experience of Regret: What, When, and Why," *Psychological Review*, 1995, 102 (2), 379–395.
- , —, and **Serena Chen**, "Commission, Omission, and Dissonance Reduction: Coping with Regret in the "Monty Hall" problem," *Personality and Social Psychology Bulletin*, 1995, 21 (2), 182–190.
- Gleicher, Faith, Kathryn A. Kost, Sara M. Baker, Alan J. Strathman, Steven A. Richman, and Steven J. Sherman**, "The Role of Counterfactual Thinking in Judgments of Affect," *Personality and Social Psychology Bulletin*, 1990, 16 (2), 284–295.
- Gollier, Christian and Alexander Muermann**, "Optimal Choice and Beliefs with Ex Ante Savoring and Ex Post Disappointment," *Management Science*, 2010, 56 (8), 1272–1284.
- Granberg, Donald and Thad A. Brown**, "The Monty Hall Dilemma," *Personality and Social Psychology Bulletin*, 1995, 21 (7), 711–723.
- Grether, David M.**, "Bayes Rule as a Descriptive Model: The Representativeness Heuristic," *Quarterly Journal of Economics*, 1980, 95 (3), 537–557.
- Herweg, Fabian and Klaus M. Schmidt**, "Loss Aversion and Inefficient Renegotiation," *Review of Economic Studies*, 2015, 82 (1), 297–332.
- and **Konrad Mierendorff**, "Uncertain Demand, Consumer Loss Aversion, and Flat-Rate Tariffs," *Journal of the European Economic Association*, 2013, 11 (2), 399–432.
- Hoffman, Mitchell**, "How is Information Valued? Evidence from Framed Field Experiments," 2014. Working Paper.
- and **Stephen V. Burks**, "Training Contracts, Worker Overconfidence, and the Return from Firm-Sponsored General Training," 2015. Working Paper.
- Hong, Harrison, Terence Lim, and Jeremy C. Stein**, "Bad News Travels Slowly: Size, Analyst Coverage, and the Profitability of Momentum Strategies," *Journal of Finance*, 2000, 55 (1), 265–295.
- Ikenberry, David L. and Sundaresh Ramnath**, "Underreaction to Self-Selected News Events: The Case of Stock Splits," *Review of Financial Studies*, 2002, 15 (2), 489–526.
- , **Graeme Rankine, and Earl K. Stice**, "What Do Stock Splits Really Signal?," *Journal of Financial and Quantitative Analysis*, 1996, 31 (3), 357–375.
- , **Josef Lakonishok, and Theo Vermaelen**, "Market Underreaction to Open Market Share Repurchases," *Journal of Financial Economics*, 1995, 39 (2), 181–208.
- Jegadeesh, Narasimhan and Sheridan Titman**, "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency," *Journal of Finance*, 1993, 48 (1), 65–91.
- Kahneman, Daniel and Amos Tversky**, "Subjective Probability: A Judgement of Representativeness," *Cognitive Psychology*, 1972, 3, 430–454.
- and —, "On the Psychology of Prediction," *Psychological Review*, 1973, 80 (4), 237–251.
- and —, "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 1979, 47 (2), 263–292.
- and —, "The Psychology of Preferences," *Scientific American*, 1982, 246 (1), 160–173.
- and **Dale T. Miller**, "Norm Theory: Comparing Reality to its Alternatives," *Psychological Review*, 1986, 93 (2), 136–153.
- , **Jack L. Knetsch, and Richard H. Thaler**, "Experimental Tests of the Endowment Effect and the Coase Theorem," *Journal of Political Economy*, 1990, 98 (6), 1325–1348.
- Karlsson, Niklas, George Loewenstein, and Duane Seppi**, "The Ostrich Effect: Selective Attention to Information," *Journal of Risk and Uncertainty*, 2009, 38 (2), 95–115.
- Kőszegi, Botond**, "Health anxiety and patient behavior," *Journal of Health Economics*, 2003, 22 (6), 1073–1084.
- , "Ego Utility, Overconfidence, and Task Choice," *Journal of the European Economic Association*,

- 2006, 4 (4), 673–707.
- , “Emotional Agency,” *Quarterly Journal of Economics*, 2006, 121 (1), 121–155.
- **and Matthew Rabin**, “A Model of Reference-Dependent Preferences,” *Quarterly Journal of Economics*, 2006, 121 (4), 1133–1165.
- **and** —, “Reference-Dependent Risk Attitudes,” *American Economic Review*, 2007, 97 (4), 1047–1073.
- **and** —, “Reference-Dependent Consumption Plans,” *American Economic Review*, 2009, 99 (3), 909–936.
- Klibanoff, Peter, Owen Lamont, and Thierry A. Wizman**, “Investor Reaction to Salient News in Closed-End Country Funds,” *Journal of Finance*, 1998, 53 (2), 673–699.
- Kuhnen, Camelia M. and Brian Knutson**, “The Influence of Affect on Beliefs, Preferences, and Financial Decisions,” *Journal of Financial and Quantitative Analysis*, 2011, 46 (03), 605–626.
- , **Sarah Rudorf, and Bernd Weber**, “Stock Ownership and Learning from Financial Information,” 2015. Working Paper.
- Landier, Augustin**, “Wishful Thinking: A Model of Optimal Reality Denial,” 2000. Working Paper.
- Landman, Janet**, “Regret and Elation Following Action and Inaction: Affective Responses to Positive Versus Negative Outcomes,” *Personality and Social Psychology Bulletin*, 1987, 13 (4), 524–536.
- Loomes, Graham and Robert Sugden**, “Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty,” *Economic Journal*, 1982, 92 (368), 805–824.
- Loughran, Tim and Jay R. Ritter**, “The New Issues Puzzle,” *Journal of Finance*, 1995, 50 (1), 23–51.
- Macara, Rosario**, “Dynamic Beliefs,” *Games and Economic Behavior*, 2014, 87, 1–18.
- Malmendier, Ulrike and Timothy Taylor**, “On the Verges of Overconfidence,” *Journal of Economic Perspectives*, 2015, 29 (4), 3–8.
- Mayraz, Guy**, “Priors and Desires: A Model of Optimism, Pessimism, and Cognitive Dissonance,” 2014. Working Paper.
- Michaely, Roni, Richard H. Thaler, and Kent L. Womack**, “Price Reactions to Dividend Initiations and Omissions: Overreaction or Drift?,” *Journal of Finance*, 1995, pp. 573–608.
- Möbius, Markus M., Muriel Niederle, Paul Niehaus, and Tanya S. Rosenblat**, “Managing Self-Confidence,” 2014. Working Paper.
- Moutsiana, Christina, Caroline J. Charpentier, Neil Garrett, Michael X. Cohen, and Tali Sharot**, “Human Frontal-Subcortical Circuit and Asymmetric Belief Updating,” *Journal of Neuroscience*, 2015, 35 (42), 14077–14085.
- , **Neil Garrett, Richard C. Clarke, R. Beau Lotto, Sarah-Jayne Blakemore, and Tali Sharot**, “Human Development of the Ability to Learn from Bad News,” *Proceedings of the National Academy of Sciences*, 2013, 110 (41), 16396–16401.
- Mullainathan, Sendhil**, “A Memory-Based Model of Bounded Rationality,” *Quarterly Journal of Economics*, 2002, 117 (3), 735–774.
- Nicolle, Antoinette, Stephen M. Fleming, Dominik R. Bach, Jon Driver, and Raymond J. Dolan**, “A Regret-Induced Status Quo Bias,” *Journal of Neuroscience*, 2011, 31 (9), 3320–3327.
- Oster, Emily, Ira Shoulson, and E. Ray Dorsey**, “Optimal Expectations and Limited Medical Testing: Evidence from Huntington Disease,” *American Economic Review*, 2013, 103 (2), 804–830.
- , —, **and** —, “Optimal Expectations and Limited Medical Testing: Evidence from Huntington Disease: Corrigendum,” *American Economic Review*, 2016, 106 (6), 1562–1565.
- Palacios-Huerta, Ignacio**, “Learning to Open Monty Hall’s Doors,” *Experimental Economics*, 2003, 6 (3), 235–251.
- Petrocelli, John V. and Anna K. Harris**, “Learning Inhibition in the Monty Hall Problem: The Role of Dysfunctional Counterfactual Prescriptions,” *Personality and Social Psychology Bulletin*, 2011, 37 (10), 1297–1311.
- Rabin, Matthew**, “Inference by Believers in the Law of Small Numbers,” *Quarterly Journal of*

Economics, 2002, 117 (3), 775–816.

— , “An Approach to Incorporating Psychology into Economics,” *American Economic Review: Papers & Proceedings*, 2013, 103 (3), 617–622.

— **and Dimitri Vayanos**, “The Gambler’s and Hot-Hand Fallacies: Theory and Applications,” *Review of Economic Studies*, 2010, 77 (2), 730–778.

— **and Joel L. Schrag**, “First Impressions Matter: A Model of Confirmatory Bias,” *Quarterly Journal of Economics*, 1999, 114 (1), 37–82.

Ritov, Ilana and Jonathan Baron, “Reluctance to Vaccinate: Omission Bias and Ambiguity,” *Journal of Behavioral Decision Making*, 1990, 3 (4), 263–277.

Samuelson, William and Richard Zeckhauser, “Status Quo Bias in Decision Making,” *Journal of Risk and Uncertainty*, 1988, 1 (1), 7–59.

Schwartzstein, Joshua, “Selective Attention and Learning,” *Journal of the European Economic Association*, 2014, 12 (6), 1423–1452.

Sharot, Tali, Christoph W. Korn, and Raymond J. Dolan, “How Unrealistic Optimism is Maintained in the Face of Reality,” *Nature Neuroscience*, 2011, 14 (11), 1475–1479.

— , **Ryota Kanai, David Marston, Christoph W. Korn, Geraint Rees, and Raymond J. Dolan**, “Selectively Altering Belief Formation in the Human Brain,” *Proceedings of the National Academy of Sciences*, 2012, 109 (42), 17058–17062.

Shleifer, Andrei, *Inefficient Markets: An Introduction to Behavioral Finance*, Oxford University Press, 2000.

Sloman, Steven A., Philip M. Fernbach, and York Hagmayer, “Self-deception requires vagueness,” *Cognition*, 2010, 115 (2), 268–281.

Slovic, Paul and Sarah Lichtenstein, “Comparison of Bayesian and Regression Approaches to the Study of Information Processing in Judgment,” *Organizational Behavior and Human Performance*, 1971, 6, 649–744.

Spiegler, Ran, “On Two Points of View Regarding Revealed Preferences and Behavioral Economics,” in Andrew Caplin and Andrew Schotter, eds., *The Foundations of Positive and Normative Economics: A Handbook*, Oxford, UK: Oxford University Press, 2008, pp. 95–115.

Spieß, D. Katherine and John Affleck-Graves, “Underperformance in Long-Run Stock Returns Following Seasoned Equity Offerings,” *Journal of Financial Economics*, 1995, 38 (3), 243–267.

Thaler, Richard, “Toward a Positive Theory of Consumer Choice,” *Journal of Economic Behavior & Organization*, 1980, 1 (1), 39–60.

Tversky, Amos and Daniel Kahneman, “Belief in the Law of Small Numbers,” *Psychological Bulletin*, 1971, 76 (2), 105–110.

— **and** — , “Judgment under Uncertainty: Heuristics and Biases,” *Science*, 1974, 185 (4157), 1124–1131.

— **and** — , “Loss Aversion in Riskless Choice: A Reference-Dependent Model,” *Quarterly Journal of Economics*, 1991, 106 (4), 1039–1061.

Weinberg, Bruce A., “A Model of Overconfidence,” *Pacific Economic Review*, 2009, 14 (4), 502–515.

Wiswall, Matthew and Basit Zafar, “How Do College Students Respond to Public Information about Earnings?,” *Journal of Human Capital*, 2015, 9 (2), 117–169.

Yariv, Leeat, “I’ll See It When I Believe It – A Simple Model of Cognitive Consistency,” 2005. Working Paper.