Attention and Regret*

Martin Vaeth†

October 24, 2024

Abstract

This paper explains regret as an optimal self-control mechanism to motivate attention and thus improve decision-making. The model endogenizes the optimal emotions as incentives for an agent that overweights the cost of attention, for example, due to temptation or present bias. If the realized state is ex-post observable, the model provides a foundation for regret theory, including disproportionate aversion to large regrets. Advancing regret theory, the model explains why regret is stronger than rejoicing and why it is stronger in simpler decision problems. If the realized state is imperfectly observable, the model predicts a combination of regret and disappointment.

[Regret] may also depend on the extent to which the individual blames himself for his original decision. [...] The neglect of this dimension of regret - although a useful simplifying assumption for many problems - is a serious obstacle to the development and generalisation of regret theory. - Sugden, 1986

I   Introduction

Decisions under uncertainty can trigger regret and rejoicing, when the decision maker learns that another action would have done better or worse. These emotions are often considered futile "crying over spilled milk," and their anticipation has been shown to distort choices, posing a puzzling question: Why do such emotions exist if only the actual outcomes should matter to an agent, particularly from an evolutionary perspective?[1]

Although the anticipation of regret and rejoicing can distort decision-making, we show these emotions can nevertheless improve the quality of decisions if we take into account *attention*. When facing uncertainty, individuals can improve decision-making by paying costly attention, which we understand broadly as encompassing both external information acquisition and internal information processing, such as deliberation, planning, or memory search. The anticipation of regret and rejoicing can motivate individuals to pay more attention to avoid regret and seek rejoicing. Such an additional motive to pay attention is beneficial, if the agent is biased toward insufficient attention, due to the temptation to avoid attention effort or due to present bias, because benefits are only realized after the cost of attention is incurred.

We show regret and rejoicing are indeed the optimal emotions to counteract a bias for insufficient attention. We take a principal-agent approach (e.g., Robson, 2001; Samuelson and Swinkels, 2006; Rayo and Becker, 2007) to study the optimal emotions that evolution, whether natural or cultural, would have endowed us with for that purpose. We allow the principal to choose *any* emotional payoffs as a function of the chosen action and ex-post feedback about the realized state of the world. Under complete feedback, the model endogenizes regret theory, including disproportionate aversion to large regrets. Intuitively, regret can be seen as a form of *self-reproach*, punishing the agent for shirking on attention, and rejoicing as self-congratulation for the converse. Coherent with this interpretation, our model predicts regret is stronger following a mistake in simple decision problems than in harder ones. The model also explains why the optimal mix of self-punishments and rewards is tilted toward the former, that is, why regret is stronger than rejoicing. Finally, we show that under incomplete feedback, disappointment and elation are additionally used as second-best motivators. Disappointment and elation also arise from counterfactual comparisons, but in contrast to regret and rejoicing, they concern counterfactual outcomes had another *state* realized.

In our model, the principal, whether nature or a parent, delegates decision-making including attention (or information acquisition, depending on the interpretation) to an agent. Both parties care about the attention cost as well as the decision outcome, but the agent overweights the attention

---

[1]Bell (1982) and Loomes and Sugden (1982) study the behavioral implications of regret and rejoicing theoretically. Zeelenberg (1999a) reviews evidence that their anticipation affects choices.

cost, which we refer to as the agent's bias. The principal can incentivize the agent by endowing them with emotions, which enter the agent's utility. Emotions can be chosen as a function of the chosen action and feedback on the realized state, which captures two constraints (discussed in more detail in section II). First, emotions do not depend on the amount of attention paid, which we assume is difficult to verify, inducing a moral-hazard concern. Second, we allow for an uncertain bias of the agent, on which emotions cannot condition, and study the resulting effects of adverse selection. Such uncertainty may be due to the stochastic nature of temptation or to heterogeneity in present bias. To investigate the incentive problem pertaining not just to the amount but also the type of obtained information, we allow the agent to pay attention by acquiring *any* signal structure subject to an entropy-based cost as in the rational-inattention literature (Sims, 2003).

Adopting a principal-agent approach not only allows the use of methods from mechanism design but also facilitates the broader applicability of our results. On a wider scale, we demonstrate the utility of counterfactual benchmarks for incentivizing biased agents to pay attention. In section II and Appendix A, we comment on the relevance of our findings for contracting with experts and relate them to relative performance evaluation.

The principal-agent problem we study is complicated by the presence of both adverse selection and moral hazard, alongside the high dimensionality of the agent's strategy due to flexible information acquisition. To render the problem tractable, we impose two symmetry assumptions. First, we assume the delegated decision problem is symmetric, that is, invariant under relabeling of actions. Second, we confine ourselves to mechanisms where the emotional payoffs preserve this symmetry of the decision problem.

We find the optimal emotions are *regret* and *rejoicing*. Both arise from counterfactual thinking, because they result from comparisons with what might have happened had one chosen differently. Formally, let $v(a, \omega)$ denote the payoff under action a and state ω. Regret or rejoicing experienced in state ω and under action a is a function of the vector of payoff differences $v(a, \omega) - v(a', \omega)$ to alternative actions a' under the realized state ω. The anticipation of regret and rejoicing increases the agent's incentive to pay attention, because the agent aims not only to achieve a better decision outcome but also to avoid regret and seek rejoicing. Although various emotions can motivate attention by increasing the stakes of the decision problem, regret and rejoicing *emerge as optimal* in our model, because how much attention the agent pays to a certain state is determined by the utility differences between actions in that state. Under an unknown bias, how much the principal wishes to amplify these utility differences depends on the preexisting payoff differences, exclusive of emotions. Regret and rejoicing, as defined above, achieve precisely this amplification by being functions of the payoff differences to counterfactual actions.

Focusing on decision problems with two actions, which have received significant attention in the regret literature, we show our main results. The optimal mechanism *endogenizes* classical regret theory (Bell, 1982; Loomes and Sugden, 1982), including disproportionate aversion to large regrets. This aversion, crucial to regret theory, leads to common deviations from expected utility theory, such as the Allais paradox, the common ratio effect, and simultaneous gambling and insurance. Our

model predicts a more general version of this property than assumed in aforementioned studies. However, we demonstrate our more general version is equivalent to the conjunction of the above-mentioned deviations from expected utility. Restricting ourselves to symmetric decision problems, we show that although these choice distortions do not come into play, attention is distorted. We show disproportionate aversion to large regrets skews attention towards extreme events, that is, states where the payoff difference between the actions is greater.

Our model advances regret theories in several dimensions. First, it addresses the self-recrimination aspect of regret, which is inadequately captured in classical regret theory, as argued by one of its pioneers (Sugden, 1986). One implication is that a realistic theory of regret should take into account whether one could have known another action would turn out better (ibid.). In our framework, regret and rejoicing can be seen as punishments and rewards for paying (insufficient) attention. Indeed, the model predicts that when finding out the right action is less costly, mistakes are more severely punished by regret. Second, existing models typically focus solely on regret, overlooking rejoicing, or are agnostic about the relative intensity of both emotions (e.g., Bell, 1982; Loomes and Sugden, 1982; Quiggin, 1994; Sarver, 2008). By contrast, our model predicts both regret *and* rejoicing, while confirming the intuition that regret is stronger than rejoicing.[2] Third, our model extends regret theory to more than two actions, which has proven difficult for classical regret theory (Loomes and Sugden, 1982; Loomes and Sugden, 1987). A frequently used model for general choice sets considers only regret to the ex-post best action (Quiggin, 1994). Our model makes the psychologically plausible prediction that regret and rejoicing are affected by all counterfactual actions. In particular, under a known bias, each counterfactual action affects regret and rejoicing in proportion to the probability with which the action should have been taken according to the first-best attention strategy (Theorem 3). This probability provides a notion of choiceworthiness and addresses the other main point in Sugden (1986): self-blame is affected by another action to the "extent to which an action is a serious candidate for choice." Finally, we extend regret theory to explicitly incorporate the role of feedback, which we turn to next.

In an extension of the model, emotions can be chosen as any function of the chosen action and *incomplete feedback* about the state. With complete feedback, the principal can identify not only the realized payoff but also counterfactual payoffs from alternative actions, which regret relies on. However, complete feedback is not always plausible. For instance, if the agent is selecting between different economic projects, only the return of chosen project and general market conditions may be observable ex post, but not the return of alternative projects. To model such situations, we assume the principal can condition emotions only on the payoff of the chosen action, and on a partitional signal of the realized state, such as market conditions.

Under incomplete feedback and an unknown bias, all emotional payoffs have to be determined jointly, making the solution analytically intractable. Therefore, we simplify by assuming a known bias, which can be perfectly offset by a combination of regret and rejoicing with another set of

---

[2]Although it turns out that the distinction between regret and rejoicing does not matter for choice behavior (e.g. Loomes and Sugden, 1982), it has implications for welfare comparisons and menu choice (Sarver, 2008).

emotions, namely, *disappointment* and *elation*. Our prediction aligns with empirical evidence, that regret and rejoicing are predominant under complete feedback, whereas disappointment and elation gain prominence when only the realized payoff is known (Zeelenberg, 1999a; Camille et al., 2004). To our knowledge, we are the first to incorporate both sets of emotions as a function of the feedback structure. Furthermore, under a known bias, the result is robust to discarding the above-mentioned symmetry assumptions and to generalizing the entropy-based attention cost to any posterior-separable cost function.

This paper also has implications for how to evaluate the emotions of regret and disappointment. Knowing their benefits and costs is useful, especially when we can influence these emotions through socialization or therapy. In this context, our work contributes to a broader debate on whether emotions are adaptive (Keltner and Gross, 1999). More specifically, Camille et al. (2004) find patients who do not experience regret as the result of damage to their orbitofrontal cortex make worse decisions under risk, in line with the hypothesis that regret is an adaptive emotion to improve decision-making (e.g., Janis and Mann, 1977; Zeelenberg, 1999b; O'Connor et al., 2012). The existing well-developed hypothesis as to *how* regret and disappointment improve decision-making is that they serve as heuristics for learning. When an agent faces a decision under uncertainty repeatedly, reinforcement learning with regret (Foster and Vohra, 1999) or disappointment and elation (Schultz et al., 1997) have desirable asymptotic properties. By contrast, we maintain the assumption of Bayesian learning and show the *anticipation* of regret and disappointment can serve as optimal motivators for costly learning. Although we believe these two approaches can be complementary, our approach is able to explain more properties of regret, such as disproportionate aversion to large regrets and the relevance of the difficulty of the decision problem as well as of feedback.

Related Literature   Regret and disappointment are the primary emotions studied in the context of decision-making under uncertainty (Zeelenberg et al., 2000; Loewenstein and Lerner, 2003). Our paper endogenizes regret theory, formalized by Bell (1982) and Loomes and Sugden (1982) for decision problems with two actions. Bleichrodt and Wakker (2015) review evidence supporting regret theory and its applications to financial and health-related decisions. The concept of disappointment was introduced into economics by Bell (1985), who posits loss aversion around the expected payoff similar to prospect theory (Kahneman and Tversky, 1979), and Loomes and Sugden (1986), who assume disappointment and elation are convex. Gul (1991) axiomatizes a theory of disappointment similar to Bell's, though the reference point here is the expected utility inclusive of emotions. Kőszegi and Rabin (2006) propose a model of gain-loss utility around a reference point, which is stochastic and distributed according to recent beliefs. Loss aversion around a reference point is *not optimal* to motivate attention. However, in Vaeth (2022), we show that under a sufficiently bimodal distribution of the bias, such loss aversion around a reference point or goal can be optimal to induce effort at the decision-making stage, as opposed to the attention stage. In contrast to the above-mentioned papers, we endogenize regret and disappointment from their hypothesized

function as internal self-control mechanisms.

At a broader level, our work contributes to a large body of research, initiated by Demski and Sappington (1987), that explores how a principal should incentivize an expert to pay attention, or, in the terminology of that literature, to acquire information. Our model distinguishes itself by incorporating flexible information acquisition and both moral hazard, due to the acquired information not being contractible, and adverse selection, due to the agent's bias being private information of the agent.[3]

Although Osband (1989) also addresses both moral hazard and adverse selection, his model is confined to a more specific setup. The principal employs an expert to predict a variable under a quadratic loss, with the expert drawing independent observations at a cost unknown to the principal. The author concludes the optimal contract is quadratic in prediction, linear in realization, and that self-screening offers little value. Our model imposes fewer constraints on the decision problem and allows the agent to acquire any signal structure at a posterior-separable cost, as seen in more recent contributions to this field (Rappoport and Somma, 2017; Carroll, 2019; Lindbeck and Weibull, 2020; Yoder, 2022; Whitmeyer and Zhang, 2022). In this regard, we leverage advances in rational inattention and information design, particularly those of the posterior approach (Caplin and Dean, 2013; Caplin et al., 2022), using concavification methods from Bayesian persuasion (Kamenica and Gentzkow, 2011), as well as the state-dependent choice-probabilities approach (Matějka and McKay, 2015).

Yoder (2022) studies adverse selection and flexible information acquisition in a model with two states, but he assumes the information-acquisition strategy, or at least the realized signal, is contractible. By contrast, our model only considers the chosen action and realized state as contractible. We show the principal cannot make use of self-screening, and hence, our solution methods and the resulting distortions differ: in his model, there is no distortion at the top, whereas in our case, information acquisition is skewed toward extreme states for all types.

Another difference from the above-mentioned papers is that in our model, the agent has a *direct interest* in the decision, albeit being biased towards insufficient information acquisition. This feature is shared with Szalay (2005), but in his paper, the bias is known and the main result concerns the optimal mechanism when the principal has no access to transfers. He shows that, absent transfers, the principal can incentivize the agent by restricting the choice set and allowing only extreme actions.

This paper is part of a larger project that relates the optimal mechanism in the face of an uncertain bias to non-standard preferences. In Vaeth (2022), we investigate the optimal emotions when the agent's bias does not affect attention but does affect the decision-making stage, for instance, because some actions demand more effort. We give conditions for when the induced utility function replicates loss aversion around a reference point or goal and when mental accounting is used as an incentive instrument. Our work has relevance to the literature on contract theory with

---

[3]This type of adverse selection at the outset is to be distinguished from adverse selection in the interim, resulting from the realization of information acquisition being private information of the agent. The latter type of adverse selection, which is the focus of Thereze (2022), is also present in our model.

behavioral agents (Kőszegi, 2014), specifically O'Donoghue and Rabin (2006), Lockwood (2020) and Beshears et al. (2020), who explore optimal mechanisms to disincentivize overconsumption and incentivize effort and savings, respectively, for agents with heterogeneous present bias.

The paper is organized as follows. Section II introduces and discusses the model. Section III solves the model when the principal has complete feedback about the realized state and the bias is unknown. Section IV extends the model to incomplete feedback about the state, when the bias is known. Section V concludes.

II   Model

Decision Problem   The principal (evolution, she) faces a *decision problem*, which is a tuple (Ω, μ, A, v). It comprises a finite state space Ω, a full-support prior $\mu \in \Delta(\Omega)$ on the state space, a finite set of actions A, and a utility function $v : A \times \Omega \to \mathbb{R}$, which we refer to as the payoff.[4]

Attention   The principal delegates a two-stage problem to the agent (he). First, he acquires a signal structure about the state (the attention stage) and, second, he chooses an action from A conditional on the realized signal (the decision-making stage). A signal structure is a pair (S, σ) where S is a finite signal space and $\sigma : \Omega \to \Delta(S)$ is the state-dependent distribution over signals. Following the rational-inattention approach (Sims, 2003), we assume that the agent can select *any* signal structure subject to the *mutual information* cost, defined below. It is well known (e.g. Matějka and McKay, 2015) that the agent's signal can be identified with an action recommendation, reducing the agent's two-stage problem to the selection of state-dependent choice probabilities $p : \Omega \to \Delta(A)$, written $p(a|\omega)$. The state-dependent choice probabilities p, which we also refer to as an *attention strategy*, entail an attention cost $c(p)$ that corresponds to the mutual information of the action and state,

$$c(p) := \sum_{a \in A} -p(a) \log p(a) - \sum_{\omega \in \Omega} \sum_{a \in A} -\mu(\omega) p(a|\omega) \log p(a|\omega).$$

Here, $p(a) := \sum_{\omega \in \Omega} \mu(\omega) p(a|\omega)$ is the unconditional probability of a and $\log$ refers to the natural logarithm. The first sum is the entropy of the unconditional choice probabilities, while the second is the expected entropy of the conditional choice probabilities. Intuitively, entropy quantifies the uncertainty of a distribution and mutual information measures how much knowing the state reduces the uncertainty about the action. Hence, the attention cost reflects the extent to which the action is tailored to the state.

---

[4]We assume a finite state space to circumvent measure theoretical issues. We conjecture that results carry over to infinite state spaces under the appropriate technical assumptions.

Preferences   The principal's utility V derived from an attention strategy p is

$$V(p) = -\lambda c(p) + \sum_{\omega \in \Omega} \sum_{a \in A} \mu(\omega) p(a|\omega) v(a, \omega).$$

That is, she internalizes the attention cost, weighted by λ, and derives utility $v(a, \omega)$ from action a under state ω. We interpret λ as the difficulty of the decision problem, as it scales the cost of learning about the state. This parameter will be used for performing comparative statics.

As in the literature on delegation, the agent has a bias, represented here as a discount factor β between 0 and 1 applied to the utility $v(a, \omega)$. Equivalently, this bias can be viewed as an overweighting of the attention cost by a factor of $1/\beta$. In our main interpretation, the bias captures a present bias, or the temptation to avoid the attention cost.

We depart from standard delegation models in two respects. First, the bias β is uncertain and private information of the agent, with cumulative distribution function F. This uncertainty can be due to incomplete knowledge of the agent's fixed bias as in models of self-signaling (Bénabou and Tirole, 2004) or due to the stochastic nature of the bias, as in random-indulgence models (Chatterjee and Krishna, 2009; Duflo et al., 2011; Dekel and Lipman, 2012).[5] Second, the principal can use action- and state-contingent payoff-transfers $m(a, \omega)$ to incentivize the agent. We interpret these transfers as *emotions*, which can depend on the chosen action and the realized state.[6] The agent with bias β derives utility U_β from the attention strategy p,

$$U_\beta(p) = -\lambda c(p) + \beta \sum_{\omega \in \Omega} \sum_{a \in A} \mu(\omega) p(a|\omega) \big(v(a, \omega) + m(a, \omega) \big).$$

Equivalently, the agent maximizes

$$\frac{U_\beta(p)}{\beta} = -\frac{\lambda c(p)}{\beta} + \sum_{\omega \in \Omega} \sum_{a \in A} \mu(\omega) p(a|\omega) \big(v(a, \omega) + m(a, \omega) \big).$$

Principal's Problem and Tie-Breaker   Thus far, the model can only identify optimal emotions up to a state-wise constant. Given that the agent is an expected-utility maximizer, shifting all emotional payoffs in a given state by a constant does not alter his behavior.[7] Moreover, emotions do not enter the principal's utility directly because, as usual, evolution just cares about the induced behavior. Thus, the principal is indifferent between any shifts of emotions by state-wise constants. Yet, while such constants are not relevant for the agent's behavior, they do matter for welfare comparisons.

---

[5]In the terminology Dekel and Lipman (2012), our agent is a random Strotz type (as opposed to a random Gul-Pesendorfer type).

[6]As common in many models, we assume that emotions are additively separable in utility (e.g. Bell, 1982, 1985; Loomes and Sugden, 1982, 1986). We also make this assumption because it allows the principal to implement any emotion-inclusive utility function $u(a, \omega) = v(a, \omega) + m(a, \omega)$. Up to determining a state-wise constant, this is assumption is immaterial for the model's prediction of u.

[7]This is a subtle issue. As in regret theory, given emotions, the agent is an expected-utility maximizer. However, since emotions can be a function of the choice set, the agent's choice behavior can violate expected-utility axioms.

To resolve such indifferences, we introduce two refinements. First, we impose an ex-ante budget constraint (BC). We set the budget b to zero and interpret the budget constraint as a consequence of hedonic adaptation.[8] As the agent gets used to the emotions, their effect on utility is zero in expectation.

$$\int \left(\sum_{\omega \in \Omega} \sum_{a \in A} \mu(\omega) p_\beta(a|\omega) m(a, \omega) \right) dF(\beta) = b \tag{BC}$$

Second, from the possible solutions to the principal's problem, we refine the principal's choice to the mechanism that minimizes the variance of emotions,

$$\int \left(\sum_{\omega \in \Omega} \sum_{a \in A} \mu(\omega) p_\beta(a|\omega) (m(a, \omega) - b)^2 \right) dF(\beta).$$

This variance minimization can be understood as the principal's effort to reduce a quadratic implementation cost associated with inducing emotions.[9] This cost could be due to the energy used to induce emotions in the brain or an opportunity cost from diminished capacity for other cognitive processes. To ensure that the optimal incentives for attention shape the payoff structure rather than the implementation cost, we treat the minimization of variance as a tie-breaker, rather than directly incorporating it into the principal's utility function.

Interpretation and Discussion Our model aligns with the principal-agent approach to the evolution of preferences (e.g. Robson, 2001; Samuelson and Swinkels, 2006; Rayo and Becker, 2007). This approach likens evolution to a principal that chooses preferences to incentivize fitness-maximizing behavior while operating under certain contracting constraints. In our model, these constraints include: (1) the agent is inherently biased, and (2) the degree of this bias and attention are private information of the agent.

Regarding the first constraint, we assume that evolution cannot directly mitigate the bias, such as temptation. Instead, emotions like regret and disappointment function as kludges to improve decision-making. This assumption aligns with evidence for an incremental evolution of the mind, which adds new structures while leaving old structures intact (Linden, 2007; Marcus, 2008).[10] Furthermore, this view is consistent with standard views of self-control issues as manifestations of a divided self. Dual-self models, also known as planner-doer models or dual-system models of the brain, conceptualize the organism as a far-sighted self interacting with a short-sighted self (Thaler and Shefrin, 1981; Bénabou and Pycia, 2002; Bernheim and Rangel, 2004; Fudenberg and Levine, 2006, 2012; Brocas and Carrillo, 2008). These two selfs can be identified with evolutionary newer and older brain structures, such as the prefrontal cortex and the amygdala, which aligns with neuroscientific evidence of the prefrontal cortex exerting control over other brain structures (Miller and Cohen, 2001; see also Brocas and Carrillo, 2008).[11] Compatible with this, neuroscientific

---

[8]If $b \neq 0$, the optimal emotions are simply uniformly shifted by a constant.

[9]As we show in the Appendix, most of our results apply under a more general implementation cost.

[10]Ely (2011) presents a theoretical foundation for why an organism might be kludged.

[11]While our main interpretation of the principal is evolution creating brain structures that induce regret, a dual-self

research points toward a pivotal role of the prefrontal cortex in experiencing regret (Camille et al., 2004; Coricelli et al., 2005).

Regarding the second constraint, we assume that evolution cannot condition emotions on the acquired signal structure or on the stochastic bias β. This might stem from the fact that once the outcome materializes, it is easier to recall and verify actions taken than the intensity of temptation or the amount of attention paid before decision-making, as is assumed in models of self-signalling (Bodner and Prelec, 2003; Bénabou and Tirole, 2004).

Another, possibly complementary, interpretation of the principal is as cultural evolution or parents. Emotions may be culturally transmitted, for instance, through parental praise or reprimands for children's decisions, leading to the internalization of such feedback as self-congratulatory (rejoicing) or self-blaming (regret) emotions. Under this interpretation, the first constraint above amounts to assuming that culture or parents cannot affect an underlying present bias or temptation directly. The second constraint means that the degree of the bias and attention are private information of the child and parental feedback can condition only on the feedback about the decision outcome.

Finally, our model can be applied to contractual relationships, where information acquisition is delegated to an expert and payoff-transfers are monetary instead of emotions. For example, in corporate settings, shareholders entrust CEOs with the responsibility of selecting viable projects, and in financial markets, investors delegate the choice of investment portfolios to fund managers. For risk-neutral parties, we can interpret our transfers as monetary, subject to an ex-ante budget constraint. This constraint allows the principal to deviate from the budget as long as the the *expected* transfers do not exceed it. Among the contracts that optimize the agent's behavior and comply with the ex-ante budget constraint, we assume that the principal chooses the one that minimizes the variance of the ex-post budget imbalance. Under this interpretation, the agent's bias need not denote a behavioral bias but can also capture a divergence of interests between the principal and the agent. We discuss these two interpretations further in Appendix A.

The reader who prefers a more standard economic interpretation of the model, is invited to reinterpret emotions in the following as monetary transfers.

III   Complete Feedback

In this section, we maintain the assumption that the principal can condition emotions on the chosen action a and the realized state ω. The next section analyzes the contracting problem when only a partition of the state is contractible.

We assume that the distribution of the agent's type β is non-degenerate, so there is some uncertainty about the bias of the agent. Let $\text{supp}(\beta)$ denote the support of the distribution of the agent's type. We make the assumption that $\inf(\text{supp}(\beta)) > 0$, so we can show later that finite emotions are optimal. Further, we assume for convenience that the distribution of β is discrete,

---

interpretation is also possible. Novel compared to existing dual-self models is that our principal can use emotions as rewards or penalties to incentivize the agent. An exception is papers on goal-setting (Koch and Nafziger, 2011; Hsiaw, 2013), however, they constrain the choice of a utility function to specific functional forms.

which simplifies some parts of the proof of Lemma 1 and Theorem 2.

Assumption 1. *The bias β non-degenerate, discrete distribution with* $\mathrm{supp}(\beta) \subseteq [\underline{\beta}, 1]$ *where* $\underline{\beta} > 0$.

Due to the presence of both moral hazard and adverse selection in the principal's problem, the problem is very complex in general. In particular, the agent's choice of an attention strategy is an high-dimensional problem and generally does not have analytic solutions. To render the problem tractable, we make two symmetry assumptions. First, we impose that relative to the choice set A and prior μ, the utility v is *exchangeable* (Matějka and McKay, 2015). Second, we restrict ourselves to mechanisms where emotional payoffs maintain exchangeability of the utility. Exchangeable utility captures that the decision problem invariant under relabeling the actions. It helps to characterize the agent's attention strategy, which is given by a multinomial logit formula (Matějka and McKay, 2015). This allows us to reduce the principal's problem to a state-wise optimization problem.

Definition 1 (Exchangeability). *A utility function $u : A \times \Omega \to \mathbb{R}$ is exchangeable if the distribution over payoff vectors $(u(a,\omega))_{a \in A} \in \mathbb{R}^A$ induced by prior μ is invariant under permutation of the coordinates.*

As an example, the utility function according to the payoff-matrix in Table 1 is exchangeable if $\mu(\omega_1) = \mu(\omega_2)$. That is because the payoff vector $(0,1)$ is as likely as the payoff vector $(1,0)$. If either $\mu(\omega_1) \neq \mu(\omega_2)$ or if we changed one of the payoffs, then the utility would not be exchangeable anymore.

	ω_1	ω_2
a_1	1	0
a_2	0	1

Table 1: Exemplary payoff-matrix

Assumption 2. *The utility v is exchangeable.*

This symmetry assumption is natural, for example, when before any attention the actions are indistinguishable and arguably should hold at a sufficiently ex-ante stage. Under our main interpretation, this assumption is plausible if emotions are adapted to the general attention problem that the agent faces and not to specific contexts.

For tractability, we focus on emotions m that maintain exchangeability of the decision problem. Thus, the emotions do not depend on the labeling of actions. This symmetry assumption is natural because the underlying decision problem is symmetric in the actions, so the first-best attention strategy is symmetric in the actions.[12]

---

[12]Simulations under two states suggest that inducing asymmetries through emotions cannot benefit the principal. For a quadratic posterior-separable cost under two states, we can show that the asymmetric emotions are suboptimal (proof upon request).

We study the following principal's problem P.

$$\max_{m,(p_\beta)_\beta} \int \left(-c(p_\beta) + \sum_{\omega \in \Omega} \sum_{a \in A} \mu(\omega) p_\beta(a|\omega) v(a,\omega) \right) dF(\beta) \tag{P}$$

s.t.

$$\forall \beta : p_\beta \in \arg\max_p -c(p) + \beta \sum_{\omega \in \Omega} \sum_{a \in A} \mu(\omega) p(a|\omega)\big(v(a,\omega) + m(a,\omega)\big) \tag{IC}$$

$$\int \left(\sum_{\omega \in \Omega} \sum_{a \in A} \mu(\omega) p_\beta(a|\omega) m(a,\omega) \right) dF(\beta) = 0 \tag{BC}$$

$$v + m \text{ is exchangeable.} \tag{EX}$$

Out of the solutions to P, the principal chooses the one that minimizes the variance of emotions m.

If the agent's bias β was known, there would be a simple way to offset it through emotions. Any emotions that raise the stakes of the decision problem by $1/\beta$, for example $m(a,\omega) = \left(\frac{1}{\beta} - 1\right) v(a,\omega)$, perfectly align the interest of the agent and principal, since

$$\beta\left(v(a,\omega) + \left(\frac{1}{\beta} - 1\right) v(a,\omega)\right) = v(a,\omega).$$

Thus, different types require differently strong incentives to offset their bias.

Therefore, one might wonder whether by offering a menu of emotion schemes, the principal can tailor incentives to the agent's type.[13] Even more generally, the principal could elicit a message from the agent before and after attention prior to assigning an emotion scheme, to screen the agent. Assuming that every emotion scheme has to induce an exchangeable utility, we show that under two states, the principal cannot do better than offering a single emotion scheme, so the formulation of P above is without loss.[14]

Lemma 1. *Let $|\Omega| = 2$. Under Assumption 1 and 2, any mechanism is weakly dominated by a single emotion scheme and the formulation P is without loss.*

All proofs are relegated to the Appendix. The proof shows that offering a menu makes types sort into emotion schemes to the principal's disadvantage. Ideally, the principal would want to offer stronger incentives to types with lower β than to types with higher β because the former group is more biased. However, higher types have a higher return to attention and pick stronger incentives than lower types, contrary to the principal's aim. This leads to higher types paying more attention and lower types less than under a single emotion scheme. Hence, offering a menu increases the variance of attention, decreasing the principal's utility.

First, we characterize the solution of P for decision problems containing two actions and then turn to the case of $n > 2$ actions.

---

[13] In analogy to the term "transfer schemes", we also use "emotion scheme" to refer to the function m.

[14] We conjecture that this Proposition generalizes to any (finite) state space.

A Two Actions

Let $A = \{a_1, a_2\}$. To see whether the optimal emotions can be interpreted as regret and rejoicing, we define a regret solution of the principal's problem, which is based on classical regret theory in Bell (1982) and Loomes and Sugden (1982). For a_i, $i = 1, 2$, denote by a_{-i} the other action in A, so $a_{-1} = a_2$ and $a_{-2} = a_1$.

Definition 2 (Regret Solution). *Under two actions, $A = \{a_1, a_2\}$, a solution $(m, (p_\beta)_\beta)$ to the principal's problem* P *is a regret solution if there exists a function $R : \mathbb{R} \to \mathbb{R}$ such that for all $(a_i, \omega) \in A \times \Omega$:*

$$m(a_i, \omega) = R\big(v(a_i, \omega) - v(a_{-i}, \omega)\big)$$

and

1. *$R(x) \gtreqqless 0 \Leftrightarrow x \gtreqqless 0$,*

2. *$\frac{R(x) - R(-x)}{x}$ is strictly increasing on $\mathbb{R}_{>0}$.*

Note that a regret solution does not only require that the agent's preferences can be represented by a utility function as in regret theory, but further that the emotional utility m itself can be written as the regret utility term in Loomes and Sugden (1982) and Bell (1982). That is because, like the emotional utility term in regret theory, m should be interpretable as experienced utility. The function R is called a regret-rejoicing function in regret theory and captures the emotions of regret (if $v(a_i, \omega) - v(a_{-i}, \omega) < 0$) and rejoicing (if $v(a_i, \omega) - v(a_{-i}, \omega) > 0$). Thus, in regret theory the agent cares about the outcomes of actions but also anticipates and cares about the emotional consequences, which enter the utility function as an additive component.

Property 1 states that if the actual and the counterfactual payoff are the same ($x = 0$), the agent experiences neither regret nor rejoicing ($R(x) = 0$). If the actual payoff is better than the counterfactual payoff ($x > 0$), then the emotion is positive ($R(x) > 0$) and if the counterfactual payoff is better ($x < 0$), then the emotion is negative ($R(x) < 0$).[15]

Property 2 requires that the sum of regret and rejoicing for some payoff difference is increasing more than proportionally. It is often called disproportionate aversion to large regrets and captures that large payoff differences get overweighted by regret, which leads to violations of expected-utility axioms. In fact, Property 2 *captures all the behavioral consequences of regret* for attention and decision-making. To see that, first, note that Property 2 is equivalent to requiring that

$$Q(x) := x + R(x) - R(-x)$$

satisfies $\frac{Q(x)}{x}$ is strictly increasing on $\mathbb{R}_{\geq 0}$, by virtue of $\frac{Q(x)}{x} = 1 + \frac{R(x) - R(-x)}{x}$. Further, as the agent is an expected-utility maximizer, adding a state-dependent constant to utility does not change

---

[15]Loomes and Sugden (1982) make the stronger assumption that $R(x)$ is strictly increasing in x, that is they assume the better the actual compared to the counterfactual outcome, the better the emotion. The model of this paper does predict this property for negative x, but for large enough $x > 0$, $R(x)$ can be decreasing. In the context of the model, this can be interpreted as saying that for large utility differences, there is nothing to congratulate oneself for because any type would have paid enough attention to choose the right action with stakes that large.

incentives for attention or decision-making. Hence, behavior is determined by the *stakes* in each state, that is the state-wise utility difference between the two actions. Denoting the *emotion-inclusive utility*

$$u(a,\omega) := v(a,\omega) + m(a,\omega),$$

Q is exactly the function that maps the stakes without emotions,

$$\Delta v(a_i,\omega) := v(a_i,\omega) - v(a_{-i},\omega),$$

into the stakes including emotions,

$$\Delta u(a_i,\omega) := u(a_i,\omega) - u(a_{-i},\omega)$$

by

$$\Delta u(a_i,\omega) = \Big(v(a_i,\omega) + R\big(v(a_i,\omega) - v(a_{-i},\omega)\big)\Big) - \Big(v(a_{-i},\omega) + R\big(v(a_{-i},\omega) - v(a_i,\omega)\big)\Big)$$
$$= \Delta v(a_i,\omega) + R\big(\Delta v(a_i,\omega)\big) - R\big(-\Delta v(a_i,\omega)\big)$$
$$= Q(\Delta v(a_i,\omega)).$$

Property 2 is more general than the assumption in Bell (1982) and Loomes and Sugden (1982) that Q is strictly convex, which implies that $\frac{Q(x)}{x}$ is strictly increasing on $\mathbb{R}_{>0}$.

Despite Property 1 and 2 being less restrictive than the assumptions in Bell (1982) and Loomes and Sugden (1982), we show in the Appendix E, Proposition 1, that our weaker assumptions are still sufficient for many predictions of regret theory, notably the Allais paradox, the common-ratio effect, the isolation effect, and simultaneous gambling and insurance. As we show, Property 2 is in fact equivalent to the conjunction of these effects in the context of regret theory.

The following theorem shows that principal's problem indeed has a regret solution. This is remarkable, since one could expect that the optimal emotions under some outcome depend on the whole decision problem. That is they could depend on the the payoff v of every outcome and on the prior. However, they follow the simple regret formula, which means that emotions depend only on the difference of the realized payoff to the payoff had one taken the other action, and they do so following a regret-rejoicing function, which does not depend on the state nor on the prior. However, the regret-rejoicing function depends on λ and on the distribution of β, which we use for comparative statics below.

Theorem 1. *Under a binary decision problem and Assumptions 1 and 2, the principal's problem* P *has a regret solution. The regret-rejoicing function R of such a solution is unique almost everywhere on* $\mathbb{R}$.

The intuition for why regret and rejoicing are optimal is as follows. One can show that under exchangeability, the principal's problem separates by state, so the optimal emotions depend only on the payoffs in the realized state. In any state, the agent's stochastic choice probabilities are

13

given by a multinomial logit function (Matějka and McKay, 2015):

$$p_\beta(a_i|\omega) = \frac{e^{\frac{\beta}{\lambda}\Delta u(a_i,\omega)}}{1 + e^{\frac{\beta}{\lambda}\Delta u(a_i,\omega)}} \tag{1}$$

This illustrates that the optimal attention strategy is such that the probability with which the agent picks the right action in some state depends on the (emotion-inclusive) stakes in that state. The higher the stakes, the more attention the agent pays to make sure that he takes the right action under this state. What regret and rejoicing do is exactly to increase the stakes: regret penalizes the wrong action and rejoicing rewards the right action, state by state. This explains Property 1.

The solution endogenously satisfies Property 2, or disproportionate aversion to large regrets, by the following. As shown above, Property 2 is equivalent to $Q(x)/x$ being strictly increasing on $\mathbb{R}_{>0}$, where Q maps, in each state, the stakes without emotions Δv into the emotion-inclusive stakes $\Delta u = Q(\Delta v)$. If β was known, the optimal solution would be

$$Q(\Delta v) = \frac{\Delta v}{\beta},$$

since it perfectly offsets the agent's discounting of Δv. Under unknown β, the optimal Q increases the stakes in some state to offset a *weighted average* of the bias. The average is taken with respect to the distribution of types weighted by how much these types' choice probabilities are *sensitive* to the stakes. Formally, the optimal $Q(\Delta v)$ is implicitly characterized by the ratio of v and an endogenous weighted average of the bias,

$$Q(\Delta v) = \frac{\Delta v}{\frac{\int p'_\beta(Q(\Delta v))\beta dF(\beta)}{\int p'_\beta(Q(\Delta v))dF(\beta)}}, \tag{2}$$

where

$$p'_\beta(Q(\Delta v)) = \frac{d}{dQ(\Delta v)}\frac{e^{\frac{\beta}{\lambda}Q(\Delta v)}}{1 + e^{\frac{\beta}{\lambda}Q(\Delta v)}},$$

For higher stakes, higher types are comparatively less sensitive to the stakes because they already take the right action with probability close to 1. Thus, the emotions are more targeted to the low types, which require proportionally higher emotions to offset their bias. This leads to disproportionate aversion to large regrets.

In general, equation (2) can have multiple solutions. We show that the solutions are increasing in the strong set order and thus wherever Q is non-unique it has an upward jump. The non-decreasing function Q can have only countably many jump discontinuities, from which we show that R is unique up to at most countably many points.

Finally, the principal's problem can have solutions that are not regret solutions. When (2) has multiple solutions and there are duplicate states, she could pick from the solution set as a function of the state. Then, emotions are not a state-independent function of the set of payoff-differences anymore violating our regret solution definition. However, the regret solution is arguably the

simplest type of solution because the principal can choose emotions only as a function of the vector of state-wise payoffs. In fact, she can employ the same function R *independent of the decision problem.*

Turning to the behavior induced by the optimal contract, it is useful to recall the distinction between the attention stage (choice of a signal structure) and the decision-making stage (choice of an action conditional on the signal). Corollary 1 shows that the principal only distorts behavior in the former stage.

Corollary 1. *The agent's behavior at the decision-making stage is not distorted. The agent's attention is skewed towards more extreme events.*

A common theme in the principal-expert literature is that the principal might want to distort the expert's incentives at the decision-making stage to incentivize more information acquisition, as already pointed out by Demski and Sappington (1987). In the current model, that is not necessary. Because of the symmetry implied by exchangeability, under any posterior that any type of the agent will optimally acquire, the action he takes is the same as the principal's preferred action, as we show in the Appendix. However, under other posteriors, regret and rejoicing can distort behavior, as is known from regret theory (Bell, 1982; Loomes and Sugden, 1982).

Regarding attention, Property 2 implies that the agent pays relatively more attention to more extreme states compared to the first-best. We measure extremeness of a state ω by the stakes $|v(a_1, \omega) - v(a_2, \omega)| = |\Delta v(a_1, \omega)|$. How much attention the type β agent pays to some state ω can be measured by the log-likelihood ratio of the signal under ω. Because the signal is identical to the action, equation (1) implies that the log-likelihood ratio equals

$$\log\left(\frac{p_\beta(a_1|\omega)}{p_\beta(a_2|\omega)}\right) = \log\left(e^{\frac{\beta}{\lambda}Q(\Delta v(a_1, \omega))}\right) = \frac{\beta}{\lambda}Q\big(\Delta v(a_1, \omega)\big).$$

By contrast, the log-likelihood ratio with respect to the first-best attention strategy p^* is

$$\log\left(\frac{p^*(a_1|\omega)}{p^*(a_2|\omega)}\right) = \log\left(e^{\frac{1}{\lambda}\Delta v(a_1, \omega)}\right) = \frac{1}{\lambda}\Delta v(a_1, \omega).$$

The ratio of those two log-likelihood ratios, whether for a fixed type β or in expectation over β, grows with the extremeness of the state, since $Q\big(\Delta v(a_1, \omega)\big)/\Delta v(a_1, \omega)$ is increasing as $|\Delta v(a_1, \omega)|$ increases by Property 2. Thus, the model predicts that the agent pays relatively more attention to more extreme events.

The prediction of disproportionate attention to states with a high payoff difference is similar in spirit to models of salience or focusing based on contrast (Bordalo et al., 2012; Kőszegi and Szeidl, 2013). Salience effects or focusing are often understood as a feature of automatic bottom-up attention instead of deliberate top-down attention, such as rational inattention (Bordalo et al., 2022). In contrast, we show that such focusing effects are not only consistent with a model of rational (in)attention, but that they are in fact optimal distortions to motivate rational inattention.[16]

---

[16]Lieder et al. (2018) offer a different explanation for overrepresentation of extreme events based on the resource-

The model makes two additional predictions regarding emotions that we relate to regret theory. First, while not imposed in standard regret theory, it is widely assumed that regret is stronger than rejoicing (e.g. Larrick and Boles, 1995; Humphrey, 2004), for which there is evidence based on self-reported emotions (Beattie et al., 1994; Mellers et al., 1999). Moreover, models of regret often do not include rejoicing at all (Quiggin, 1994; Sarver, 2008), suggesting that regret is considered more prominent than rejoicing.

Corollary 2. *Regret is stronger than rejoicing, that is,* $\forall x > 0 : |R(-x)| > |R(x)|$.

This result stems from the principal's motive to minimize the quadratic cost of inducing emotions.[17] Rejoicing is costly for the principal when the agent takes the right action, and regret is costly for the principal when the agent takes the wrong action. By (1), the agent will, in any state, take the right action with higher probability than the wrong action. Thus, the principal finds it cheaper to use regret than rejoicing. This motive is present in the analysis of optimal moral systems by Kaplow and Shavell (2007), who assume an inculcation cost for guilt and virtue. Further, it echoes similar motives elsewhere in game theory as expressed by Schelling (1980), "a promise is costly when it succeeds, and a threat is costly when it fails."[18]

Second, in contrast to existing models of regret, our model endogenizes the optimal mechanism as a function of the parameter λ, which can be interpreted as the difficulty of the decision problem. The higher λ the more costly it is to learn about the state and find the right action. Sugden (1986) argues that a short-coming of standard regret theory is that it does not incorporate that regret depends on how easy it was that one could have known the right action. He relates this to what he calls the self-recrimination aspect of regret. The following corollary shows that our model predicts stronger regret if it was cheaper to find out the optimal choice.

Corollary 3. *A smaller* λ *leads to larger regret from taking the wrong action, that is higher* $|R(-x)|$ *for all* $x > 0$. *Further, it leads to higher emotion-inclusive stakes, that is higher* $Q(x)$ *for all* $x > 0$.

Two mechanisms contribute to this effect. First, a smaller λ makes the types with lower β comparatively more responsive to changes in stakes, which increases $Q(x)$ and second, a smaller λ makes it more likely that the agent takes the right action, which makes regret comparatively cheaper than rejoicing for the principal.

Corollary 3 can also be related to inter-personal differences. The parameter λ not only captures properties of the decision problem but also of the agent, namely how costly it is for them to process information. A small λ corresponds to decision-makers for whom it is relatively cheap to process information or, equivalently, who have a greater information processing capacity, for example due to a smaller cognitive load. The model predicts that they will experience stronger regret when

---

rationality approach.

[17]As we argue in the Appendix A, the refinement of selecting the mechanism that minimizes a quadratic cost of inducing emotions is equivalent to the budget constraint and the minimization-of-variance refinement. We prove Corollary 2 for all costs of inducing emotions that are symmetric around zero, with quadratic costs as a special case.

[18]Note that while penalties are stronger than rewards, this does not mean that the emotions are negative in expectation. In fact, they cancel out in expectation by the ex-ante budget constraint.

taking the wrong action. This can help explain why Schwartz et al. (2002) and Iyengar et al. (2006) find that "maximizers" are less satisfied with their choices than "satisficers" despite the first group making better choices on average.

B More than Two Actions

As stressed in the literature, it is not trivial how to generalize regret theory to more than two actions (Loomes and Sugden, 1982, Loomes and Sugden, 1987, Quiggin, 1994). We define a *general regret solution* based on the utility representation in Loomes and Sugden (1987) and Sugden (1993), which to our knowledge accommodates all proposals. They define a representation of preferences where the utility depends on the realized outcome and on the set of counterfactual outcomes under the same state. They do not impose additional assumptions that would correspond to generalizations of Properties 1 and 2. In fact, it is not trivial how these would generalize to more than two actions. This representation is more general than another widely-used representation due to Quiggin (1994), according to which regret depends on the difference of the actual to the highest attainable payoff under the realized state. As we will show, that form of regret is in general not optimal in our model.

Definition 3 (General Regret Solution). *Under a finite choice set A, a solution $(m, (p_\beta)_\beta)$ to the principal's problem P is a general regret solution if there exists a function R such that for all $(a, \omega) \in A \times \Omega$:*

$$m(a, \omega) = R\big(\{v(a, \omega) - v(a', \omega)\}_{a' \in A \setminus \{a\}}\big),$$

and for all $a, a' \in A, \omega \in \Omega$:

$$v(a, \omega) > v(a', \omega) \Leftrightarrow v(a, \omega) + m(a, \omega) > v(a', \omega) + m(a', \omega).$$

The definition states that emotions are a state-independent function of the *set* of payoff differences between the actual and the counterfactual payoffs under the realized state. In particular, a permutation of the counterfactual payoffs does not change the emotion as well as shifting all payoffs under some state by a constant. Moreover, adding emotions preserves the strict ordering of payoffs, state-wise. In the Appendix C, we comment on how our representation relates exactly to the one in Loomes and Sugden (1987).

Finally, assume that the principal has strict preferences over actions, given the state. This can be interpreted as a genericity assumption or as the state being defined in a sufficiently fine-grained way. We use this assumption to prove that state-wise permutations of payoffs translate to corresponding permutations of the optimal emotions.

Assumption 3. *For all $\omega \in \Omega$, for all $a, a' \in A$: $(a \neq a' \rightarrow v(a, \omega) \neq v(a', \omega))$.*

Under these assumptions, we prove the following theorem.

Theorem 2. *Under finite A and Assumptions 1 to 3, the principal's problem P has a general regret solution.*

The main part of the proof is to show that the principal's problem separates by state and that state-wise optimal emotions exist, a result which we also use for proving Theorem 1. Moreover, we show that state-wise optimal emotions are invariant under state-wise permutations or shifts of payoffs, and preserve strict orders of payoffs.[19]

Our model does not only predict that the optimal emotions are a general regret solution. Given λ and the distribution of β, the model can be used to compute the precise optimal emotional payoffs. While optimal emotions could, again, depend on the whole decision problem, they boil down to a state-wise $|A|$-variable optimization problem. That is, the optimal emotions are a function of the vector of payoffs in the realized state. They are given by the variance-of-emotions-minimizing solution to

$$\max_{u(\cdot,\omega):A\to\mathbb{R}} \int \left(\lambda H\big(p_\beta(u(\omega))\big) + \sum_{a\in A} p_\beta(a|u(\omega))v(a,\omega) \right) dF(\beta),$$

where $p_\beta(a|u(\omega))$ is given by the multinomial logit formula

$$p_\beta(a|u(\omega)) = \frac{e^{\frac{\beta}{\lambda}u(a,\omega)}}{\sum_{a'\in A} e^{\frac{\beta}{\lambda}u(a',\omega)}},$$

and $H(p_\beta(u(\omega)))$ is the entropy of the probability vector $p_\beta(u(\omega)) = (p_\beta(a|u(\omega)))_{a\in A}$. Unfortunately, the problem does not allow for a simple analytical solution or for comparative statics as in binary decision problems. However, once the the problem is solved for a vector of payoffs, it applies to payoff-equivalent states in *any decision problem*. Further, when the bias is known, the next section shows that there is a simple analytic and intuititve solution (Theorem 3).

IV   Incomplete Feedback

So far, we have assumed that the principal can choose emotions as an arbitrary function of the action and realized state. In many decision contexts, however, ex post only the payoff is revealed but not the full state. For example, if the agent pursues some economic project, the realized return might be observable but not the return had he pursued another project. In particular, when only the realized payoff $v(a,\omega)$ is learned but not the state ω, the principal cannot in general infer the payoffs of counterfactual actions. Thus, she cannot implement the general regret solution from the previous section. This section generalizes the principal's problem to account for incomplete feedback under the simplifying assumption of a known bias.

As noted earlier, under known bias it is possible to implement the first-best attention strategy by offering a contract that scales up the payoff-differences by $1/\beta$. However, to find the emotion scheme m with minimal variance, we need to determine *all* emotion schemes that implement the optimal attention strategy. Below, we argue how this exercise teaches us something about the

---

[19]Solutions that are not general regret solutions exist for the same reason as in the case with two actions. Further, one can show analogously to the binary action case that the principal distorts behavior only at the attention stage and not at the decision-making stage.

18

nature of the optimal mechanism when the first-best attention strategy cannot be implemented due to an uncertain bias.

Ex-post feedback also plays an important role in the study of regret and disappointment. While classical regret theory (Bell, 1982; Loomes and Sugden, 1982) does not take into account feedback, several experiments have documented that regret affects decisions more if feedback is expected (Zeelenberg, 1999a). To accommodate for that evidence, Humphrey (2004) introduces a theory of feedback-conditional regret which employs two regret-rejoicing functions, one for complete feedback and one for incomplete feedback. Gabillon (2020) proposes a regret theory for arbitrary feedback structures, where the regret-rejoicing function is applied to the difference between the actual payoff and the highest counterfactual payoff, given the posterior beliefs after feedback. Evidence from Camille et al. (2004) based on self-reported emotions, skin conductance responses, and choices suggests, however, that under partial feedback (only the payoff of the chosen action is learned) primarily disappointment and elation are experienced, while under complete feedback mainly regret and rejoicing are experienced. We are not aware of theories that combine regret and disappointment depending on the feedback structure.

To incorporate feedback into our model, we assume that the principal can choose emotions only as a function of the feedback that is received ex post. Formally, we assume the emotion scheme m is measurable with respect to a feedback structure $\mathcal{F} : A \times \Omega \to Z$ where Z is a message space, that is,

$$\forall a, a' \in A, \omega, \omega' \in \Omega : \mathcal{F}(a, \omega) = \mathcal{F}(a', \omega') \to m(a, \omega) = m(a', \omega').$$

We consider partitional feedback structures that reveal the payoff $v(a, \omega)$, the action a, and a partition $\mathcal{P}$ of the state space.[20] In one extreme case, the feedback structure reveals the state fully (the finest partition; complete feedback) and in the other extreme, only the payoff is revealed (the coarsest partition; partial feedback). In intermediate cases, some information about the state is revealed, such as general market conditions, while more specific information, such as how each particular economic project would have done, is not revealed.

In some decision-problems, however, the agent can back out more information about the state from the material payoff and the action. To rule out such cases, we make an additional assumption that the message space is a "connected" on each partition. To formulate this condition, we endow the message space, Z, with a symmetric relation $\sim$, which makes Z an undirected graph where $\sim$ stands for the edges of the graph. We define $z \sim z'$ if and only if there exists $\omega \in \Omega$ and $a, a' \in A$ such that $z = \mathcal{F}(a, \omega)$ and $z' = \mathcal{F}(a', \omega)$. Thus, two messages are directly connected by an edge if they are observed under the same state. Because the state-wise payoff-differences determine the incentives for attention, this is the relevant notion of connectedness that ensures that the emotions

---

[20]A partition of the state space is a set of mutually exclusive and exhaustive subsets $\mathcal{P} = \{P_1, ..., P_n\}$ of Ω. In the Appendix, we show that the resulting emotions are essentially of the same form, when we allow for other deterministic feedback structures that reveal at least the payoff. Further, none of the results below would change if the action was not revealed.

of the two messages have to be determined jointly. More specifically, the following condition ensure that the set of messages consistent with a partition cannot be divided into disjoint subsets such that emotions can be chosen independently on each subset.

Assumption 4 (Partitional Feedback Structure). *There is a partition $\mathcal{P}$ of Ω such that*

$$\mathcal{F}(a, \omega) = (v(a, \omega), a, \mathcal{P}(\omega))$$

where $\mathcal{P}(\omega)$ is the partition cell that contains ω. We define Z as the image of $\mathcal{F}$ in $\mathbb{R} \times A \times \mathcal{P}$.

For each partition cell $P \in \mathcal{P}$, the subset of the message space that is consistent with this partition cell, $\{(v, a, P) \in Z | v \in \mathbb{R}, a \in A\}$, is a connected subgraph of Z.

For simplicity, we analyze the model under a known bias $\beta \in (0, 1)$. Thus, we consider a pure moral hazard problem without an adverse selection component.

Assumption 5. *The distribution of $\beta \in (0, 1)$ is degenerate.*

Under known β, we can generalize the model in two dimensions. First, we drop the exchangeability assumption on v and the restriction to mechanisms with exchangeable $v + m$. Second, we generalize the entropy-based attention cost to any posterior-separable cost function. To be coherent with the notation above, we continue to define the attention cost in terms of the state-dependent stochastic choice function.[21] For each state-dependent stochastic choice function $p : \Omega \to \Delta(A)$, let $\tau(p) \in \Delta(\Delta(\Omega))$ denote the distribution over posteriors $\gamma \in \Delta(\Omega)$ induced by p via Bayesian updating.

Definition 4 (Posterior-Separable attention cost). *The attention cost c is posterior-separable if*

$$c(p) = \mathbb{E}_{\tau(p)}[H(\mu) - H(\gamma)]$$

where $H : \Delta(\Omega) \to \mathbb{R}$ is strictly concave and differentiable on $\Delta(\Omega)$.

Posterior-separable attention costs arise naturally from several foundations, including information theory (Sims, 2003), sequential sampling (Morris and Strack, 2019; Hébert and Woodford, 2023), and constant marginal cost of experimentation (Pomatto et al., 2023). The function H can be understood as a measure of uncertainty of a distribution, with entropy as an example. Thus, under a posterior-separable cost, the agent incurs a cost proportional to the expected reduction of the uncertainty of their belief. That H is strictly concave ensures that more information (in the Blackwell sense) is more costly.

Assumption 6. *The attention cost $c(p)$ is posterior separable.*

---

[21] This is without loss if the function H below is strictly concave. As is well-known, if we defined the attention cost as a function of the acquired signal structure, we can without loss identify the signal with the action under strictly concave H.

Given incomplete feedback and the absence of adverse selection, the new principal's problem $\mathrm{P}^{\mathrm{inc}}$ is to pick the solution with minimal variance of emotions to:

$$\max_{m,p} -c(p) + \sum_{\omega \in \Omega} \sum_{a \in A} \mu(\omega)p(a|\omega)v(a,\omega) \qquad (\mathrm{P}^{\mathrm{inc}})$$

s.t.

$$p \in \arg\max_{\tilde{p}} -c(\tilde{p}) + \beta \sum_{\omega \in \Omega} \sum_{a \in A} \mu(\omega)\tilde{p}(a|\omega)\big(v(a,\omega) + m(a,\omega)\big) \qquad (\mathrm{IC})$$

$$\sum_{\omega \in \Omega} \sum_{a \in A} \mu(\omega)p(a|\omega)m(a,\omega) = 0 \qquad (\mathrm{BC})$$

$$\forall a, a' \in A, \omega, \omega' \in \Omega : \mathcal{F}(a,\omega) = \mathcal{F}(a',\omega') \rightarrow m(a,\omega) = m(a',\omega') \qquad (\mathrm{F})$$

We introduce the following class of solutions, which we interpret below as a combination of regret and disappointment. To simplify the exposure, we assume that the first-best attention strategy p, which maximizes the principal's objective, is unique. Further, we define the consideration set $S \subseteq A$ as the set of actions that have positive probability according to the first-best attention strategy p, formally $S = \{a \in A | \sum_{\omega \in \Omega} \mu(\omega)p(a|\omega) > 0\}$.

Definition 5 (Linear Regret-Disappointment Solution). *A solution (m,p) to the principal's problem $\mathrm{P}^{\mathrm{inc}}$ is a linear regret-disappointment solution if for all $a \in S, \omega \in \Omega$:*

$$m(a,\omega) = \left(\frac{1}{\beta} - 1\right)\big(v(a,\omega) - v(\mathcal{P}(\omega))\big) \qquad (3)$$

where $v(\mathcal{P}(\omega))$ is the expected payoff conditional on partition element $\mathcal{P}(\omega)$ with respect to the first-best attention strategy p,

$$v(\mathcal{P}(\omega)) = \frac{1}{\mu(\mathcal{P}(\omega))} \sum_{\omega' \in \mathcal{P}(\omega)} \sum_{a \in S} \mu(\omega')p(a|\omega')v(a,\omega').$$

Such a solution has the natural interpretation as setting for each partition cell a reference point or benchmark $v(\mathcal{P}(\omega))$, which corresponds to the expected payoff the agent should achieve – according to the first-best attention strategy – conditional on that partition cell.[22] Note that in contrast to other models of reference points (Kahneman and Tversky, 1979; Gul, 1991; Kőszegi and Rabin, 2006), the optimal gain-loss utility here is linear.

It is instructive to compare (3) to the regret theory by Gabillon (2020), where regret depends on the difference between the realized payoff and the highest payoff that could have been achieved through some action given the ex-post feedback about the state. As in their proposal, the reference point is updated on feedback about the state.[23] However, in our model, the reference point is not

---

[22]In the Appendix, we show that for non-quadratic costs of emotions, an analogous solution with a different reference point v is optimal.

[23]To be precise, in (3), the reference point is updated only on feedback about the state that is action-independent, namely what is learned through the partition of the state space. It is not updated on what is learned from the realized payoff $v(a,\omega)$.

the highest attainable counterfactual payoff but instead the *expected* payoff conditional on ex-post feedback. Because this expected payoff is also affected by the payoffs of the chosen action *had different states realized*, one can interpret the emotional payoff as a combination of regret and disappointment. Next, we make this intuition more clear.

Another, mathematically equivalent, formulation of (3) relates the emotional payoff to regret and disappointment. Let $\gamma : A \to \Delta(\Omega)$ map each action to the posterior associated with it via Bayesian updating, which we write $\gamma(\omega|a)$. We can define, *conditional on partition cell* $\mathcal{P}(\omega)$, the probability, $p(a|\mathcal{P}(\omega))$, and the expected payoff of, $v(a|\mathcal{P}(\omega))$, of a,

$$p(a|\mathcal{P}(\omega)) := \sum_{\omega' \in \mathcal{P}(\omega)} \frac{\mu(\omega')}{\mu(\mathcal{P}(\omega))} p(a|\omega')$$

$$v(a|\mathcal{P}(\omega)) := \sum_{\omega' \in \mathcal{P}(\omega)} \frac{\gamma(\omega|a)}{\gamma(\mathcal{P}(\omega)|a)} v(a,\omega).$$

Given this notation, we can rearrange terms and rewrite equation (3) as

$$
\begin{aligned}
m(a,\omega) = &\sum_{a \neq a'} p(a'|\mathcal{P}(\omega)) \left(\frac{1}{\beta} - 1 \right) \big(v(a,\omega) - v(a'|\mathcal{P}(\omega)) \big) \\
&+ p(a|\mathcal{P}(\omega)) \left(\frac{1}{\beta} - 1 \right) \big(v(a,\omega) - v(a|\mathcal{P}(\omega)) \big).
\end{aligned}
\tag{4}
$$

The term in the first line can be interpreted as regret and rejoicing. The emotion depends on the weighted difference between the actual payoff and the expected payoffs of all other actions a', conditional on partition cell $\mathcal{P}(\omega)$. The expected payoff of each action is computed with respect to the first-best attention strategy. The regret-like comparisons to different actions are weighted by the probabilities of the respective actions, again conditional on the partition element $\mathcal{P}(\omega)$.

This is in contrast to a widely used theory of regret where the comparison is with respect to the ex-post optimal action (Quiggin, 1994; Sarver, 2008), which allows for regret only and no rejoicing. This proposal is arguably extreme when the agent has a large choice set and does not seem psychologically plausible when the ex-post optimal action was not even considered. Moreover, in some relevant applications as in finance, under no short-sale constraints the ex-post optimal return is infinite (Qin, 2020). In contrast, our prediction corresponds to the proposal of Loomes and Sugden (1982) to generalize regret theory to multiple actions based on such a weighted sum of comparisons to *all* other actions, not just the ex-post optimal one. They highlight the importance of providing a theory of the appropriate action weights. Sugden (1986) writes that

> "Any plausible theory of such weights must, I think, take some account of the extent to which an action is a serious candidate for choice or, to put it slightly differently, of the extent to which the individual could sensibly blame himself for not having chosen it."

In our contexts, the action weights are naturally provided by the state-dependent choice probabilities under the first-best attention strategy. These choice probabilities capture with what probability

the agent should have taken some action under the realized state. This implies that regret and rejoicing do not depend on actions outside of the consideration set, which are taken with probability zero. This captures that as Sugden (1986) writes, only actions that are "serious candidates" or "real options" should be sources of regret and rejoicing. While regret and rejoicing are linear here and thus predict no violations of expected utility axioms from choice, linear regret still has effects on menu choice (Sarver, 2008).

The term in the second line of (4) can be interpreted as disappointment and elation. The emotion depends on the different between the payoff $v(a, \omega)$ and the expected payoff of the chosen action. The expected payoff of the chosen action is computed with respect to the posterior that the agent holds when taking the action. However, in contrast to standard models of disappointment and elation (Bell, 1985; Loomes and Sugden, 1986; Gul, 1991), the expected payoff is conditioned on what is learned about the state. The weight on disappointment and elation is the probability of taking the actual action conditional on the partition cell.

Theorem 3. *Under Assumptions 4 to 6, any solution of* $\mathrm{P}^{\mathrm{inc}}$ *is a linear regret-disappointment solution and induces the first-best attention strategy.*

Because we assumed that the bias is known, the optimal mechanism can perfectly offset it by scaling up the utility differences by $1/\beta$. Using a result in Whitmeyer and Zhang (2022), which builds on the Lagrangian Lemma of Caplin et al. (2022), we show that such inflating of the utility differences is also *necessary*, in order to implement the first-best attention strategy. The remaining goal of the principal is to reduce the variance of emotions, while scaling up the utility differences. This goal is constrained by the feedback structure. Under complete feedback, the second line of (4) is zero, and only regret is used. However, for incomplete feedback structures, both lines of (4) are generally non-zero and accordingly the optimal emotions can be interpreted as a combination of regret and disappointment.

Disappointment and elation are used because, under incomplete feedback, the choice of all emotional utilities consistent with some partition element of states is interrelated (Appendix, Lemma 10). Thus, all payoffs under that partition element, including the payoffs of the chosen action *had another state realized*, affect the optimal emotions under some feedback signal.[24]

Thus, in our model, disappointment and elation can be thought of as *second-best* motivators for attention when complete feedback is not available. They are second-best because they lead to a higher variance of emotions than using regret and rejoicing.[25]

One might be surprised that disappointment and elation motivate more attention because they depend on the realization of the state, which the agent cannot affect. By contrast, regret and rejoicing seem to depend more directly on the agent's action, which perhaps more intuitively incentivizes attention. However, although the agent cannot choose the unconditional distribution of states, he can implicitly choose the distribution of states *conditional* on an action by paying

---

[24]This statement holds also under unknown bias.

[25]Under an unknown bias, as we have shown under complete feedback, regret and rejoicing are not only optimal because they minimize the variance of emotions but also because they allow tailoring the incentives to the state.

attention. Indeed, this is the perspective of attention taken by the posterior approach, in which the agent acquires posteriors over states conditional on each taken action (Caplin and Dean, 2013). This goes to say that, like regret and rejoicing, the emotions of disappointment and elation penalize bad outcomes and reward good outcomes and thus raise the return to attention. As argued above, regret and rejoicing raise the return to attention in the optimal way but under constraints through the feedback structure, disappointment and elation have a role to play.

V Conclusion

While counterfactual thinking is often dismissed as irrational (e.g. Howard, 1992)[26] and seen as distortions to rational decision-making, the present model shows that emotions based on counterfactual reasoning are actually optimal to motivate better decision-making. This paper uses a rational-inattention model, but any model that trades off decision-making costs with better decision outcomes would arguably predict a useful role for regret and disappointment to amplify the stakes of the decision problem. Although the model allows for any emotions as a function of ex-post feedback, simple and well-known emotions arise endogenously. Under complete feedback about the state of the world, regret and rejoicing are optimal. In a binary decision problem, these emotions feature disproportionate aversion to large regrets and imply that attention is skewed towards extreme states. When feedback about the state is incomplete, disappointment and elation supplement regret and rejoicing as a form of second-best emotions. Extending standard regret theory, the model predicts stronger regret the easier one could have known the right action, consistent with the psychology of a self-blaming aspect of regret.

Many open questions remain. Our analysis is for tractability and conciseness restricted by symmetry assumptions. Can the principal profit from asymmetric incentive schemes in symmetric decision problems? What are the optimal incentives when the underlying decision problem is asymmetric? When delegating asymmetric decision problems, motivating attention may come only at the cost of distorting decision-making. Analyzing the optimal incentives could help shed further light on the resulting trade-off and thereby unite themes from the current paper, which shows that regret is optimal to motivate attention absent distortions of choice, and studies in regret theory, which focus on how regret distorts decision-making.

On a broader scale, our research underscores parallels between internal psychological mechanisms, such as counterfactual comparisons, and more traditional principal-agent problems. We believe that such a perspective has the potential to profit both behavioral economics and mechanism design. Behavioral economics might glean insights from mechanism design to better understand the function and benefits of non-standard preferences. Similarly, mechanism design can turn to human psychology as a source of innovative solutions for agency problems, which evolution has already solved.

---

[26]However, see Bradley and Stefánsson (2017) for a defense of counterfactual comparisons as rational.

References

Andreoni, J., Harbaugh, W., & Vesterlund, L. (2003). The Carrot or the Stick: Rewards, Punishments, and Cooperation. *American Economic Review*, *93*(3), 893–902.

Aumann, R. J. (1997). *Rationality and bounded rationality*. Springer.

Beattie, J., Baron, J., Hershey, J. C., & Spranca, M. D. (1994). Psychological determinants of decision attitude. *Journal of Behavioral Decision Making*, *7*(2), 129–144.

Bell, D. E. (1982). Regret in Decision Making under Uncertainty. *Operations Research*, *30*(5), 961–981.

Bell, D. E. (1985). Disappointment in decision making under uncertainty. *Operations research*, *33*(1), 1–27.

Bénabou, R., & Pycia, M. (2002). Dynamic inconsistency and self-control: A planner–doer interpretation. *Economics Letters*, *77*(3), 419–424.

Bénabou, R., & Tirole, J. (2004). Willpower and personal rules. *Journal of Political Economy*, *112*(4), 848–886.

Bernheim, B. D., & Rangel, A. (2004). Addiction and cue-triggered decision processes. *American economic review*, *94*(5), 1558–1590.

Beshears, J., Choi, J. J., Clayton, C., Harris, C., Laibson, D., & Madrian, B. C. (2020). *Optimal Illiquidity* (Working Paper No. 27459). National Bureau of Economic Research.

Bleichrodt, H., & Wakker, P. P. (2015). Regret theory: A bold alternative to the alternatives. *The Economic Journal*, *125*(583), 493–532.

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

Bordalo, P., Gennaioli, N., & Shleifer, A. (2012). Salience theory of choice under risk. *The Quarterly journal of economics*, *127*(3), 1243–1285.

Bordalo, P., Gennaioli, N., & Shleifer, A. (2022). Salience. *Annual Review of Economics*, *14*, 521–544.

Bradley, R., & Stefánsson, H. O. (2017). Counterfactual Desirability. *The British Journal for the Philosophy of Science*, *68*(2), 485–533.

Brocas, I., & Carrillo, J. D. (2008). The brain as a hierarchical organization. *American Economic Review*, *98*(4), 1312–46.

Camille, N., Coricelli, G., Sallet, J., Pradat-Diehl, P., Duhamel, J.-R., & Sirigu, A. (2004). The Involvement of the Orbitofrontal Cortex in the Experience of Regret. *Science*, *304*(5674), 1167–1170.

Caplin, A., & Dean, M. (2013). *Behavioral implications of rational inattention with shannon entropy* (tech. rep.). National Bureau of Economic Research.

Caplin, A., Dean, M., & Leahy, J. (2022). Rationally Inattentive Behavior: Characterizing and Generalizing Shannon Entropy. *Journal of Political Economy*, *130*(6), 1676–1715.

Carroll, G. (2019). Robust incentives for information acquisition. *Journal of Economic Theory*, *181*, 382–420.

Chatterjee, K., & Krishna, R. V. (2009). A "Dual Self" Representation for Stochastic Temptation. *American Economic Journal: Microeconomics*, *1*(2), 148–67.

Coricelli, G., Critchley, H. D., Joffily, M., O'Doherty, J. P., Sirigu, A., & Dolan, R. J. (2005). Regret and its avoidance: A neuroimaging study of choice behavior. *Nature Neuroscience*, *8*(9), 1255–1262.

Dekel, E., & Lipman, B. L. (2012). Costly Self-Control and Random Self-Indulgence. *Econometrica*, *80*(3), 1271–1302.

Demski, J. S., & Sappington, D. E. (1987). Delegated expertise. *Journal of Accounting Research*, 68–89.

Duflo, E., Kremer, M., & Robinson, J. (2011). Nudging farmers to use fertilizer: Theory and experimental evidence from Kenya. *American economic review*, *101*(6), 2350–90.

Edmans, A., Gabaix, X., & Jenter, D. (2017). Executive compensation: A survey of theory and evidence. *The handbook of the economics of corporate governance*, *1*, 383–539.

Ely, J. C. (2011). Kludged. *American Economic Journal: Microeconomics*, *3*(3), 210–231.

Evans, R. B., Gomez, J.-P., Ma, L., & Tang, Y. (2023). Peer versus Pure Benchmarks in the Compensation of Mutual Fund Managers.

Foster, D. P., & Vohra, R. (1999). Regret in the on-line decision problem. *Games and Economic Behavior*, *29*(1-2), 7–35.

Fudenberg, D., & Levine, D. K. (2006). A dual-self model of impulse control. *American economic review*, *96*(5), 1449–1476.

Fudenberg, D., & Levine, D. K. (2012). Timing and self-control. *Econometrica*, *80*(1), 1–42.

Gabillon, E. (2020). When choosing is painful: Anticipated regret and psychological opportunity cost. *Journal of Economic Behavior & Organization*, *178*, 644–659.

Gul, F. (1991). A theory of disappointment aversion. *Econometrica: Journal of the Econometric Society*, 667–686.

Halevy, Y., & Feltkamp, V. (2005). A Bayesian approach to uncertainty aversion. *The Review of Economic Studies*, *72*(2), 449–466.

Hébert, B., & Woodford, M. (2023). Rational inattention when decisions take time. *Journal of Economic Theory*, *208*, 105612.

Heiner, R. A. (1983). The origin of predictable behavior. *The American economic review*, *73*(4), 560–595.

Holmström, B. (1979). Moral Hazard and Observability. *The Bell Journal of Economics*, *10*(1), 74–91.

Howard, R. A. (1992). In Praise of the Old Time Religion. In W. Edwards (Ed.), *Utility Theories: Measurements and Applications* (pp. 27–55). Springer Netherlands.

Hsiaw, A. (2013). Goal-setting and self-control. *Journal of Economic Theory*, *148*(2), 601–626.

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

Iyengar, S. S., Wells, R. E., & Schwartz, B. (2006). Doing Better but Feeling Worse: Looking for the "Best" Job Undermines Satisfaction. *Psychological Science, 17*(2), 143–150.

Janis, I. L., & Mann, L. (1977). *Decision making: A psychological analysis of conflict, choice, and commitment.* Free press.

Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica, 47*(2), 263–292.

Kamenica, E., & Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review, 101*(6), 2590–2615.

Kaplow, L., & Shavell, S. (2007). Moral rules, the moral sentiments, and behavior: Toward a theory of an optimal moral system. *Journal of Political Economy, 115*(3), 494–514.

Kaur, S., Kremer, M., & Mullainathan, S. (2015). Self-Control at Work. *Journal of Political Economy, 123*(6), 1227–1277.

Keltner, D., & Gross, J. J. (1999). Functional accounts of emotions. *Cognition & Emotion, 13*(5), 467–480.

Koch, A. K., & Nafziger, J. (2011). Self-regulation through goal setting. *Scandinavian Journal of Economics, 113*(1), 212–227.

Kőszegi, B. (2014). Behavioral contract theory. *Journal of Economic Literature, 52*(4), 1075–1118.

Kőszegi, B., & Rabin, M. (2006). A Model of Reference-Dependent Preferences. *The Quarterly Journal of Economics, 121*(4), 1133–1165.

Kőszegi, B., & Szeidl, A. (2013). A model of focusing in economic choice. *The Quarterly Journal of Economics, 128*(1), 53–104.

Larrick, R. P., & Boles, T. L. (1995). Avoiding Regret in Decisions with Feedback: A Negotiation Example. *Organizational Behavior and Human Decision Processes, 63*(1), 87–97.

Lieder, F., Griffiths, T. L., & Hsu, M. (2018). Overrepresentation of extreme events in decision making reflects rational use of cognitive resources. *Psychological review, 125*(1), 1.

Lindbeck, A., & Weibull, J. (2020). Delegation of investment decisions, and optimal remuneration of agents. *European Economic Review, 129*, 103559.

Linden, D. J. (2007). *The accidental mind.* Harvard University Press.

Lockwood, B. B. (2020). Optimal Income Taxation with Present Bias. *American Economic Journal: Economic Policy, 12*(4), 298–327.

Loewenstein, G., & Lerner, J. S. (2003). The role of affect in decision making. *Handbook of affective sciences* (pp. 619–642). Oxford University Press.

Loomes, G., & Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *The economic journal, 92*(368), 805–824.

Loomes, G., & Sugden, R. (1986). Disappointment and dynamic consistency in choice under uncertainty. *The Review of Economic Studies, 53*(2), 271–282.

Loomes, G., & Sugden, R. (1987). Some implications of a more general form of regret theory. *Journal of Economic Theory, 41*(2), 270–287.

Ma, L., Tang, Y., & Gómez, J.-P. (2019). Portfolio Manager Compensation in the U.S. Mutual Fund Industry. *The Journal of Finance*, *74*(2), 587–638.

Marcus, G. (2008). *Kluge: The haphazard evolution of the human mind*. Houghton Mifflin Harcourt.

Matějka, F., & McKay, A. (2015). Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model. *American Economic Review*, *105*(1), 272–298.

Mellers, B., Schwartz, A., & Ritov, I. (1999). Emotion-based choice. *Journal of Experimental Psychology: General*, *128*(3), 332.

Miller, E. K., & Cohen, J. D. (2001). An Integrative Theory of Prefrontal Cortex Function. *Annual Review of Neuroscience*, *24*(1), 167–202.

Morris, S., & Strack, P. (2019). The wald problem and the relation of sequential sampling and ex-ante information costs. *Available at SSRN 2991567*.

O'Connor, E., McCormack, T., & Feeney, A. (2012). The development of regret. *Journal of Experimental Child Psychology*, *111*(1), 120–127.

O'Donoghue, T., & Rabin, M. (2006). Optimal sin taxes. *Journal of Public Economics*, *90*(10-11), 1825–1849.

Osband, K. (1989). Optimal forecasting incentives. *Journal of Political Economy*, *97*(5), 1091–1112.

Pomatto, L., Strack, P., & Tamuz, O. (2023). The Cost of Information: The Case of Constant Marginal Costs. *American Economic Review*, *113*(5), 1360–1393.

Qin, J. (2020). Regret-based capital asset pricing model. *Journal of Banking & Finance*, *114*, 105784.

Quiggin, J. (1994). Regret theory with general choice sets. *Journal of Risk and Uncertainty*, *8*(2), 153–165.

Rappoport, D., & Somma, V. (2017). *Incentivizing Information Design* (SSRN Scholarly Paper No. ID 3001416). Social Science Research Network. Rochester, NY.

Rayo, L., & Becker, G. S. (2007). Evolutionary efficiency and happiness. *Journal of Political Economy*, *115*(2), 302–337.

Robson, A. J. (2001). The biological basis of economic behavior. *Journal of economic literature*, *39*(1), 11–33.

Samuelson, L., & Swinkels, J. M. (2006). Information, evolution and utility. *Theoretical Economics*, *1*(1), 119–142.

Sarver, T. (2008). Anticipating regret: Why fewer options may be better. *Econometrica*, *76*(2), 263–305.

Schelling, T. C. (1980). *The Strategy of Conflict: With a new Preface by the Author*. Harvard university press.

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*(5306), 1593–1599.

Schwartz, B., Ward, A., Monterosso, J., Lyubomirsky, S., White, K., & Lehman, D. R. (2002). Maximizing versus satisficing: Happiness is a matter of choice. *Journal of Personality and Social Psychology*, *83*(5), 1178–1197.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, *50*(3), 665–690.

Sugden, R. (1986). Regret, recrimination and rationality. *Recent developments in the foundations of utility and risk theory* (pp. 67–80). Springer.

Sugden, R. (1993). An Axiomatic Foundation for Regret Theory. *Journal of Economic Theory*, *60*(1), 159–180.

Szalay, D. (2005). The economics of clear advice and extreme options. *The Review of Economic Studies*, *72*(4), 1173–1198.

Thaler, R. H., & Shefrin, H. M. (1981). An Economic Theory of Self-Control. *Journal of Political Economy*, 15.

Thereze, J. (2022). Screening Costly Information.

Topkis, D. M. (1978). Minimizing a submodular function on a lattice. *Operations research*, *26*(2), 305–321.

Vaeth, M. (2022). Self-Control through Emotions: Goal-Setting, Mental Accounting, and Demand for Commitment. *Working Paper*.

Whitmeyer, M., & Zhang, K. (2022). Buying Opinions. *arXiv preprint arXiv:2202.05249*.

Yoder, N. (2022). Designing incentives for heterogeneous researchers. *Journal of Political Economy*, *130*(8), 2018–2054.

Zeelenberg, M. (1999a). Anticipated regret, expected feedback and behavioral decision making. *Journal of behavioral decision making*, *12*(2), 93–106.

Zeelenberg, M. (1999b). The use of crying over spilled milk: A note on the rationality and functionality of regret. *Philosophical Psychology*, *12*(3), 325–340.

Zeelenberg, M., Van Dijk, W. W., Manstead, A. S., & vanr de Pligt, J. (2000). On bad decisions and disconfirmed expectancies: The psychology of regret and disappointment. *Cognition & Emotion*, *14*(4), 521–541.

VI   Appendix

A   The Principal's Problem

In the main text, we assume an ex-ante budget constraint and select the solution to the principal's problem that minimizes the variance of emotions. Instead, and more general, we can assume that the principal chooses the solution that minimizes an implementation cost.

Implementation Cost   We assume that there is an implementation cost $\psi : \mathbb{R} \to \mathbb{R}$ for positive as well as negative emotions. We drop the budget constraint from the principal's problem and select the solution to the resulting problem that minimizes the expected implementation cost

$$\Psi(m) := \int \left(\sum_{\omega \in \Omega} \sum_{a \in A} \mu(\omega) p_\beta(a|\omega) \psi\big(m(a, \omega)\big) \right) dF(\beta)$$

where ψ is strictly convex, differentiable, and minimal at zero.

In our main interpretation we interpret Ψ as a cost of inducing emotions. This might be an energetic cost of inducing emotions in the brain or as an opportunity cost in terms of reduced capacity for other cognitive processes. Under the cultural-evolution interpretation, the cost can be interpreted as the cost of inculcating emotions, such as the cost of parental feedback, similar to Kaplow and Shavell (2007).

Under an interpersonal interpretation with m as monetary transfers, Ψ can be thought of as a implementation cost or as a cost of deviating from an ex-ante budget. Further, when interpreting the principal's problem in organizational contexts where monetary transfers are not used, we can interpret transfers as the reduced-form utility of reward- and punishment-actions taken by the principal. Rewards, such as praise and gifts, as well as punishments, such as making the agent perform unfavorable tasks (which require monitoring), might both be costly for the principal, as captured by ψ.[27]

Formally, this refinement is equivalent to giving the principal lexicographic preferences, where she cares about the objective $\int V(p_\beta)dF(\beta)$, and in case of indifferences, about the implementation cost, $\Psi(m)$. One could have also assumed a implementation cost that directly enters the principal's utility, which then would be

$$\int V(p_\beta)dF(\beta) - \kappa\Psi(m)$$

with $\kappa > 0$. We see the lexicographic preference choice as a simplifying assumption, that makes the model more tractable and captures the limiting case of small κ for the more realistic non-lexicographic cost. Indeed, as κ converges to 0, those preferences converge to the lexicographic ones. The case of small κ is the interesting one because only for small enough cost will emotions be chosen significantly different from zero. Further, giving the principal lexicographic preferences focuses the model on the optimal shape of emotions to induce attention rather than on the implementation cost, which depends on the application.

Under a quadratic implementation cost $\psi(x) = x^2$, the implementation-cost refinement is equivalent to the ex-ante budget constraint plus the refinement of reducing the variance of emotions. Note that the mean squared emotions equal the squared mean emotions plus the variance of emotions, as for any random variable X, $\mathbb{E}[X^2] = \mathbb{E}[X]^2 + \text{Var}[X]$. Thus, minimizing the quadratic implementation cost is equivalent to minimizing the sum of the squared expected emotions and the variance of emotions. The principal can add a constant to emotions uniformly to reduce the expected emotions to zero, without affecting the variance of emotions. Further, this does not affect the agent's behavior. Thus, minimizing the quadratic implementation cost is equivalent to an ex-ante budget constraint with $b = 0$ and the refinement of minimizing the variance of emotions.

Most of our results go through with a more general implementation cost ψ that is strictly convex, differentiable, and minimal at zero. In particular, only Theorem 3 requires a quadratic ψ (but we show that a weaker result holds under $\psi(x) = |x|^k$ with $k > 1$) and Corollary 2 requires a

---

[27]See for example Andreoni et al. (2003) for more examples of rewards and punishments in different social and institutional contexts.

symmetric ψ around zero. Otherwise, all results for hold under the general ψ.

Alternative Interpretations   As mentioned in the main text, the model can be interpreted in a more conventional way as a principal hiring an agent to acquire information. The bias of the agent need not refer to a behavioral bias but can capture a conflict of interest between principal and agent. This is the case if both principal and agent care about the information cost as well as the decision outcome but differ in how they weigh these considerations. For instance, investors or shareholders may internalize the time cost of information gathering to some degree due to the resulting decision delay or because it diverts the agent from other valuable activities. CEOs or fund managers may care about the decision outcome out of intrinsic motivation or reputational concerns. However, from the principal's perspective, the agent might overemphasize the information cost in the absence of financial incentives, potentially leading to suboptimal information acquisition. In these contexts, it is realistic to assume that information acquisition is unverifiable and only the outcome from decision-making is contractible. Information acquisition often does not produce tangible evidence and even when it does, the evidence might be hard to communicate or verify without the necessary expertise on the principal's side (Lindbeck and Weibull, 2020). Finally, we take serious the possibility that the principal is uncertain about how much the agent overweights the information cost, by incorporating adverse selection.[28]

If the bias does, however, denote a behavioral bias of the agent, then our optimal contract can be readily interpreted as the *welfare-maximizing* commitment contract for the agent, subject to an ex-ante *budget balance constraint*.[29] Furthermore, as demonstrated below, our model also applies to a *profit-maximizing* principal who hires a dynamically-inconsistent expert, subject to a *participation constraint*. Here, the principal values the decision outcome net of transfers paid to the agent, but does not care about the agent's attention cost directly. The agent, on the other hand, cares only about the transfers received net of the attention cost. Instead of a budget constraint, the principal faces an ex-ante participation constraint that guarantees the agent a certain expected utility before his bias realizes. The participation constraint effectively makes the principal internalize the agent's attention cost, rendering the problem equivalent to ours.

In these contexts, our work contribute to the understanding of compensation schemes, specifically those that assess the agent's performance relative to a peer benchmark. Relative performance evaluation is common in compensation schemes for fund managers (Ma et al., 2019; Evans et al., 2023) and, to some degree, for CEOs (Edmans et al., 2017).[30] Such benchmarks can be justified by by Holmström's (1979) informativeness criterion in standard moral-hazard settings, because subtracting exogenous common shocks to the return of the industry increases the informativeness of the performance evaluation. Our model offers a different explanation based on incentives for

---

[28]A shortcoming of this application is that our model does not feature a limited liability constraint, which binds if the budget b is low enough.

[29]At first glance, a welfare-maximizing principal would have a different objective than our principal because she would care about the agent's transfers. However, this is immaterial as transfers cancel out in expectation due to the ex-ante budget constraint.

[30]Evans et al. (2023) find that 71% of mutual funds in the US use peer-benchmarked compensation.

information acquisition. Peer comparisons can be interpreted as counterfactual comparisons, since payoffs of relevant peers inform the principal about the payoff that alternative actions would have lead to. Indeed, Theorem 3 shows that under a known bias, transfers proportional to the difference between the realized return and the expected return under the realized market conditions are optimal.

Hiring a Dynamically-Inconsistent Expert Consider a principal who hires a dynamically-inconsistent agent (expert) to acquire information.[31] The principal reaps utility $v(a, \omega)$ from action a under state ω and pays monetary transfers (wages) $m(a, \omega)$ to the agent. The agent maximizes transfers paid to them. Both have quasi-linear utility. After signing the contract, the agent faces a stochastic temptation shock to his preferences, as in random-indulgence models. Accordingly, the individual-rationality constraint needs to hold in expectation. The principal's problem is

$$\max_{m, (p_\beta)_\beta} \int \left(-c(p_\beta) + \sum_{\omega \in \Omega} \sum_{a \in A} \mu(\omega) p_\beta(a|\omega) \big(v(a, \omega) - m(a, \omega) \big) \right) dF(\beta)$$

s.t.

$$\forall \beta : \ p_\beta \in \arg\max_p -c(p) + \beta \sum_{\omega \in \Omega} \sum_{a \in A} \mu(\omega) p(a|\omega) m(a, \omega) \tag{IC}$$

$$\int \left(-c(p_\beta) + \beta \sum_{\omega \in \Omega} \sum_{a \in A} \mu(\omega) p_\beta(a|\omega) m(a, \omega) \right) dF(\beta) \geq U_0 \tag{IR}$$

If the agent was dynamically consistent ($\beta \equiv 1$), the principal would achieve the first-best by selling the firm to the agent. In the presence of a bias, the principal can be seen as maintaining control of the firm in order to provide a self-control mechanism to the agent for their mutual benefit.

By standard arguments, the (IR) constraint is binding. Inserting the (IR) constraint into the objective, we obtain the following equivalent formulation.

$$\max_{m, (p_\beta)_\beta} \int \left(-c(p_\beta) + \sum_{\omega \in \Omega} \sum_{a \in A} \mu(\omega) p_\beta(a|\omega) v(a, \omega) \right) dF(\beta) - U_0$$

s.t.

$$\forall \beta : \ p_\beta \in \arg\max_p -c(p) + \beta \sum_{\omega \in \Omega} \sum_{a \in A} \mu(\omega) p(a|\omega) m(a, \omega) \tag{IC}$$

$$\int \left(-c(p_\beta) + \beta \sum_{\omega \in \Omega} \sum_{a \in A} \mu(\omega) p_\beta(a|\omega) m(a, \omega) \right) dF(\beta) = U_0 \tag{IR'}$$

The constant $-U_0$ in the principal's objective can be ignored. Identifying m with $v + m$ in the main text, the objective and (IC) constraint are identical. The (IR') constraint can be satisfied by adding a uniform constant to emotions, so (IR') acts in effect as an ex-ante budget constraint. Adding the refinement that the principal chooses the mechanism with lowest variance of transfers, renders the

---

[31]See Kaur et al. (2015) for evidence on self-control issues at work.

problem essentially equivalent to the principal's problem in the main text.[32] The only difference is that the optimal transfers emulate the payoff v in addition to the emotions in the main formulation. That is because in this formulation the agent has no direct interest in the decision. Additionally, the optimal transfers are uniformly shifted by a constant, which depends on the outside option U_0.

B   General Mechanism

A general mechanism elicits a message from the agent *before* as well as *after* attention. Conditional on the messages, an action recommendation is sent and emotion-transfers are assigned conditional on the messages and the outcome. By the revelation principle it is sufficient to restrict attention to mechanisms, where the first message corresponds to the agent's type. We maintain the assumption that any emotion scheme induces an exchangeable decision problem for the agent. First, we analyze the case when the mechanism elicits a message only before attention. At the end of the proof, we argue that eliciting a message after attention can be reduced to a mechanism of the former case.

The principal chooses a vector of emotion schemes $m_\beta : A \times \Omega \to \mathbb{R}$ and attention strategies $p_\beta : \Omega \to \Delta(A)$, indexed by the agent's type, each, subject to incentive compatibility, budget balance, and exchangeability. The incentive-compatibility constraint below captures both the truthful reporting of one's type as well as the subsequent attention and decision-making, accounting for the possibility of double deviations. The principal's problem is

$$\max_{(m_\beta, p_\beta)_\beta} \int \left(-c(p_\beta) + \sum_{\omega \in \Omega} \sum_{a \in A} \mu(\omega) p_\beta(a|\omega) v(a, \omega) \right) dF(\beta)$$

s.t.

$$\forall \beta : \ (\beta, p_\beta) \in \arg\max_{\hat{\beta}, \hat{p}} -c(\hat{p}) + \beta \sum_{\omega \in \Omega} \sum_{a \in A} \mu(\omega) \hat{p}(a|\omega) \big(v(a, \omega) + m_{\hat{\beta}}(a, \omega) \big) \qquad \text{(IC)}$$

$$\int \left(\sum_{\omega \in \Omega} \sum_{a \in A} \mu(\omega) p_\beta(a|\omega) m_\beta(a, \omega) \right) dF(\beta) = 0 \qquad \text{(BC)}$$

$$\forall \beta : \ v + m_\beta \text{ is exchangeable.} \qquad \text{(EX)}$$

Out of the solutions $(m_\beta, p_\beta)_\beta$, the principal chooses the one that minimizes the variance of emotions,

$$\int \sum_{\omega \in \Omega} \sum_{a \in A} \mu(\omega) p_\beta(a|\omega) m_\beta(a, \omega)^2 dF(\beta).$$

Osband (1989) shows that in their model offering a menu of emotion schemes can in principle benefit the principal, but using numerical simulations he shows that these gains are very small. In contrast to Osband, in our setup, heterogeneity concerns the agent's bias instead of his cost parameter. Thus, the first-best attention strategy is the same for every agent, which the proof relies on. In Yoder (2022), the principal can effectively screen agents when the acquired posterior

---

[32]The equivalence also holds under an appropriate exchangeability constraint or when the principal chooses a menu of transfer schemes.

is contractible. We show that the moral hazard constraint implies that the principal can induce only gross value functions that are convex in the posterior. Offering a menu makes the gross value function more convex, leading to further distortions at the top and bottom. Our proof does not rely on the specific shape of mutual information and we show the proposition for any posterior-separable attention cost (see section IV for a definition) with strictly concave and differentiable H that is *symmetric in the states*.

Proof. Under exchangeability of v, if there are two states, $\Omega = \{\omega_1, \omega_2\}$, then one of the following needs to be the case. The payoff v is state-wise constant, or $\mu(\omega_1) = \mu(\omega_2) = \frac{1}{2}$ and there are two actions, $A = \{a_1, a_2\}$, with exchangeable payoffs, $v(a_1, \omega_1) = v(a_2, \omega_2)$ and $v(a_1, \omega_2) = v(a_2, \omega_1)$. If there were more than two actions and v is not state-wise constant, then more than two states are necessary to have all permutations of $(v(a, \omega))_{a \in A}$ be equiprobable.

If the payoffs are state-wise constant, then the optimal attention strategy is to pay zero attention, which is achieved at the smallest implementation cost by setting the emotions to zero for all types. Thus, all u_β are the same.

Consider, then, the other case that or $\mu(\omega_1) = \mu(\omega_2) = \frac{1}{2}$ and there are two actions, $A = \{a_1, a_2\}$, with exchangeable payoffs, $v(a_1, \omega_1) = v(a_2, \omega_2)$ and $v(a_1, \omega_2) = v(a_2, \omega_1)$. Without loss, the payoff matrix for the principal can be written as

	ω_1	ω_2
a_1	v	0
a_2	0	v

where $v > 0$.

Agent's Behavior Turning to the agent, we first analyze the agent's behavior given some $u : A \times \Omega \to \mathbb{R}$ in the menu of utility functions $(u_\beta)_\beta = (v + m_\beta)_\beta$ that the principal offers.[33] By the exchangeability and an argument analogous to the above, the payoff matrix of any such u,

	ω_1	ω_2
a_1	$u(a_1, \omega_1)$	$u(a_1, \omega_2)$
a_2	$u(a_2, \omega_1)$	$u(a_2, \omega_2)$

must satisfy one of the following.

Either u is state-wise constant, that is $u(a_1, \omega_1) = u(a_2, \omega_1)$ and $u(a_2, \omega_1) = u(a_1, \omega_1)$, or $u(a_1, \omega_1) = u(a_2, \omega_2)$ and $u(a_1, \omega_2) = u(a_2, \omega_1)$. In the first case, we can make u constant over all outcomes without changing the agent's behavior (he pays no attention in either case), so it is

---

[33] By a change of variables, the principal chooses u directly instead of m.

captured by the second case, which we analyze below. We can write the agent's payoff matrix as

	ω_1	ω_2
a_1	$\underline{u} + \Delta u$	$\underline{u}$
a_2	$\underline{u}$	$\underline{u} + \Delta u$

where $\underline{u} = u(a_1, \omega_2) = u(a_2, \omega_1)$ and $\Delta u = u(a_1, \omega_1) - u(a_2, \omega_1) = u(a_2, \omega_2) - u(a_1, \omega_2)$.

We characterize the optimal attention strategy given some chosen utility function u. Following Caplin and Dean (2013), we solve for the optimal distribution over posteriors, induced by the attention strategy, analogous to Bayesian persuasion models (Kamenica and Gentzkow, 2011). Under two states, we can identify posteriors $\gamma \in \Delta(\Omega)$ with the probability in $[0, 1]$ they assign to state ω_1, which makes this posterior approach especially attractive. The problem of agent β given u is to pick a Bayes-consistent distribution over posteriors to maximize the expectation of the *net value* function N_β (Caplin and Dean, 2013):

$$\max_{\tau \in \Delta[0,1]} \mathbb{E}_\tau[N_\beta(\gamma|u)]$$

$$\text{s.t.}$$

$$\mathbb{E}_\tau[\gamma] = \frac{1}{2}$$

Here, $N_\beta(\gamma|u)$ is the net value of posterior γ for the agent of type β given u, which is

$$N_\beta(\gamma|u) = \beta(\underline{u} + \max\{\gamma \Delta u, (1 - \gamma)\Delta u\}) + \lambda H(\gamma).$$

The first summand of the net value $N_\beta(\cdot|u)$ captures the *gross value*, or instrumental value, of attention and the second the attention cost. The gross value is continuous, convex, symmetric around $1/2$, and piecewise linear as the maximum of linear functions. The attention cost is mutual information, $H(\gamma) = \gamma \log(\gamma) + (1 - \gamma) \log(1 - \gamma)$, which is strictly concave, differentiable and symmetric around $1/2$. These are the only properties of H that we make use of for the proof.

As known from the literature on Bayesian persuasion, the solution can be obtained by concavifying the net value function $N_\beta(\cdot|u)$. As we show next, this Bayesian persuasion problem takes a particularly simple form due to symmetry.

The net value function has one global maximum if $\Delta u = 0$ and two global maxima otherwise. Restricted to $[1/2, 1]$, the net value function attains its maximum by continuity. On $[1/2, 1]$, the net value function is the sum of the linear gross value and the strictly concave attention cost, so the maximum must be unique. By symmetry and differentiability of H, we have $H'(1/2) = 0$, so the maximum is at $1/2$ only if $\Delta u = 0$. In fact, by differentiability the maximum is characterized by $-\lambda H'(\gamma) = \beta \Delta u$ if $-\lambda H'(1) > \beta \Delta u$ and $\gamma = 1$ otherwise. By symmetry of H around $1/2$, the net value function $N_\beta(\cdot|u)$ is symmetric around $1/2$, so the net value function attains another maximum at the reflection point around $1/2$.

Because the maxima are symmetric around the prior probability $1/2$, there is a unique Bayes-

consistent τ whose support is the set of global maxima. This is the unique τ that attains the highest expected posterior net value because each posterior in its support maximizes the net value function. This τ is either degenerate at $1/2$ or binary and symmetric around $1/2$. The agent selects action a_1 when the realized posterior is greater than $1/2$ and a_2 when the realized posterior is less than $1/2$, if $\Delta u > 0$ and vice versa if $\Delta u < 0$. Thus, given $v > 0$, the principal never offers an emotion-inclusive utility with $\Delta u < 0$, which would induce the agent to pay costly attention only to select the wrong action.

Thus, while the agent chooses a high-dimensional distribution over posteriors, we can characterize his choice by the one-dimensional maximizer of the net value function on $[1/2, 1]$, which corresponds to the higher posterior the agent acquires. As written above, the maximizer is characterized by $\lambda - H'(\gamma) = \beta \Delta u$ if $-\lambda H'(1) > \beta \Delta u$ and $\gamma = 1$ otherwise. Because H is strictly concave and symmetric, $-H'$ is strictly increasing on $[1/2, 1]$ with $-H'(1/2) = 0$. Thus, the maximum is strictly increasing in Δu as long as the optimal γ is less than 1 and constant afterwards.

The complete strategy of the agent consists of choosing an emotion scheme and subsequently an attention strategy. By the above, the agent's behavior can be described as choosing a net value function and subsequently the maximum of the net value function on $[1/2, 1]$. This is equivalent to choosing the maximum on $[1/2, 1]$ of the pointwise maximum $N_\beta(\cdot | (u_\beta)_\beta)$ of the collection of net value functions,

$$
\begin{aligned}
N_\beta(\gamma | (u_\beta)_\beta) &= \max_{\hat{\beta}} \{ N_\beta(\gamma, u_{\hat{\beta}}) \} \\
&= \beta \max_{\hat{\beta}} \{ \underline{u}_{\hat{\beta}} + \gamma \Delta u_{\hat{\beta}} \} + \lambda H(\gamma), \text{ for } \gamma \geq 1/2.
\end{aligned}
$$

We call $N_\beta(\gamma | (u_\beta)_\beta)$ the *effective* value function. By $\Delta u_\beta > 0$ for all β, the effective value function is supermodular on $[1/2, 1]$ in β and γ and thus the optimal γ is non-decreasing in β (Topkis, 1978). The effective value function on $[1/2, 1]$ is the sum of the maximum of increasing, linear functions and a concave attention cost. The maximum of increasing, linear functions is increasing and convex. Thus, non-decreasing γ implies a non-decreasing choice of the slope $\Delta u_{\hat{\beta}}$.

To summarize, for any type of the agent, the menu of emotion-inclusive utility functions induces an effective value function and the type's acquired signal structure can be described as picking the posterior in $[1/2, 1]$ that maximizes this net value function. We have shown that (1) a higher β implies a weakly more extreme posterior, and (2) a weakly higher slope Δu_β of the gross value function. Given the choice of an emotion-inclusive utility, (3) the agent's posterior is a strictly increasing function of the slope of the gross value function, if the posterior is below 1.

Principal's Problem Turning back to the principal, her net value function is

$$
N_P(\gamma) = \max\{\gamma v, (1 - \gamma)v\} + \lambda H(\gamma).
$$

Her value from the symmetric distribution with support $\{\gamma, 1-\gamma\}$ is just $N_P(\gamma)$ by symmetry of N_P. The principal's value function N_P is strictly concave on $[1/2, 1]$ because it is the sum of the linear gross value and the strictly concave attention cost. Strict concavity implies strict quasi-concavity, which we are using below.

Suppose for the sake of contradiction that there are types β, β' such that $u_\beta \neq u_{\beta'}$. We show that the principal strictly benefits from offering only the emotion-inclusive utility u_{β^*} of a type β^* that selects the most favorable posterior to the principal,

$$\beta^* \in \arg\max_{\beta \in \text{supp}(\beta)} N_P(\gamma_\beta).$$

Such a type β^* exists by $\text{supp}(\beta)$ being finite.

By (2), greater types than β^* select a weakly higher slope and smaller types a weakly smaller slope of the gross value function. Offering only the emotion-inclusive utility u_{β^*}, by (3) weakly increases the posterior for type less than β^*, weakly decreases the posterior for types greater than β^*, and does not change the posterior of β^*. Because the ordering of posteriors, (1), still holds due to supermodularity of the new value function, each posterior moves weakly towards γ_{β^*}. By strict quasi-concavity of the principal's value function on $[1/2, 1]$ and γ_{β^*} yielding a weakly higher value to the principal, the principal benefits weakly from the induced changes and benefits strictly from strict changes. Thus, the principal is weakly better off offering only one emotion scheme.

It remains to show that in case of indifferences, the principal cannot decrease the implementation cost by offering different emotion schemes. By the above, the principal is strictly better off if some posterior moved strictly. The only case that no posterior changed, is when the only types whose emotion-inclusive utility changed selects $\gamma = 1$ at their original and at the new emotion-inclusive utility. By construction, the new value function is below the original induced value function at $\gamma = 1$, so it entails a smaller emotional payoff after selecting the better action under the realized state. Because $\beta \leq 1$, optimality entails that the slope of the constructed gross value function is at least v. Then, the implementation cost being minimized entails positive emotions at selecting the better action under the realized state. Thus, the new emotion-inclusive utility assigning smaller emotions implies a reduction of the implementation cost.

General Mechanism   As mentioned above, the principal could use an even more general mechanism. That is, she could elicit an additional message from the agent after attention and condition the emotion scheme on both the message before and after attention. We maintain the restriction that any assigned emotion scheme m is such that $v + m$ is exchangeable. Here, we sketch how the proof can be extended to account for such mechanisms. We show that by starting with a general mechanism, we can construct a mechanism *without* a message after attention that is equally good for the principal. Then, the above proof can be applied. The argument below uses that if the agent acquires symmetric posteriors, there is no difference between eliciting a message before and after attention, as we have shown above for symmetric value functions. Using this, we show that the

principal might as well offer only one emotion scheme given the first message from the agent (thus making the second message meaningless). The argument proceeds in three steps.

First, given an incentive-compatible general mechanism, fix an arbitrary type β of the agent and consider the agent's attention problem following the first message. We can, again, by direct revelation identify the first message with the agent's type, which the agent truthfully reveals. The agent's attention problem still consists of concavifying a symmetric value function. That is because analogous to the argument above, the agent faces an effective value function, denoted by $\hat{N}_\beta$, given by the maximum of the value functions induced by the emotion schemes, which the agent can choose from after attention. Because each emotion scheme induces a symmetric value function, the maximum of the value functions is symmetric, too. This implies that the set of maximizers M_β of the value function $\hat{N}_\beta$ is symmetric around $1/2$. Any Bayes-consistent distribution with support at only the maximizers of the $\hat{N}_\beta$ is a solution to the attention problem of type β.

Second, there is a symmetric, principal's preferred solution to the attention problem of type β. The principal's value function is symmetric around $1/2$, too. Hence, the principal prefers, out of the solutions to the agent's problem, those whose support includes only posteriors that give the highest value to the principal. Formally, let $M_{P,\beta} := \arg\max_{\gamma \in M_\beta} N_P(\gamma)$, which is symmetric around $1/2$. The principal's preferred solutions to the attention problem of type β are the Bayes-consistent distributions with support in $M_{P,\beta}$. Because $M_{P,\beta}$ is symmetric and the prior is at $1/2$, there is a Bayes-consistent, binary, symmetric distribution with support in $M_{P,\beta}$.

Third, starting with a mechanism that elicits a message before and after attention, we can construct an equally good mechanism that elicits a message only before attention. As we are allowing the principal to pick her preferred solution when the agent is indifferent, for any type we can select a principal's preferred solution to his attention problem that is moreover symmetric and binary by the above. By symmetry, this solution makes any agent pick the same emotion scheme (or, send the same second message) after each acquired posterior. This solution can also be implemented by offering to every agent *only this one emotion scheme*, after the first message. Offering less choice to the agent does not not violate any other incentive-compatibility constraints. Finally, we can apply the proof above to the resulting mechanism. □

C   Proof of Theorem 2

Before proving Theorem 2, which we use below to prove Theorem 1, we comment on how our general regret representation relates to the one in Loomes and Sugden (1987).

To introduce the general regret representation in Loomes and Sugden (1987), let us define $o(a, \omega)$ as the outcome from action a under state ω. Their representation requires that the agent maximizes a utility function

$$u(a, \omega) = \Phi(o(a, \omega), \{o(a', \omega) | a \in A \setminus \{a\}\}),$$

that is a function of the actual outcome and the set of counterfactual outcomes.

Our general regret solution

$$u(a,\omega) = v(a,\omega) + R\big(\{V(a,\omega) - v(a',\omega)|a' \in A \setminus \{a\}\}\big)$$

implies their representation but makes several additional assumptions. First, the emotion-inclusive utility $u(a,\omega)$ depends on the set of state-wise *payoffs* instead of outcomes. Thus, payoff-equivalent states are treated the same way. Second, the additive separability together with the dependence of emotions on the set of payoff *differences* additionally imposes that shifting all payoffs in some state by a constant does not change the emotions. Third, we assume that emotions do not change, in each state, the strict ordering of utilities, as in Sarver (2008). These three additional assumptions seem plausible in the context of regret theory. We now turn to the proof of Theorem 2, which we show for a general implementation cost. By Appendix A, this captures as a special case the budget constraint with the refinement of minimal variance of emotions.

Proof: Let $A := \{a_1, ..., a_n\}$. Because the principal is constrained to choose emotions that induce an exchangeable emotion-inclusive utility u, the state-dependent choice probabilities of any type of the agent can be written as (Matějka and McKay, 2015)

$$p_\beta(a|\omega) = \frac{e^{\frac{\beta}{\lambda}u(a,\omega)}}{\sum_{a' \in A} e^{\frac{\beta}{\lambda}u(a',\omega)}} = \frac{e^{\frac{\beta}{\lambda}\Delta u(a,\omega)}}{1 + \sum_{a' \neq a_1} e^{\frac{\beta}{\lambda}\Delta u(a',\omega)}}. \tag{5}$$

where $\Delta u(a,\omega) := u(a,\omega) - u(a_1,\omega)$. Further, we define $u(\omega) := (u(a_1,\omega), ..., u(a_n,\omega))$ and $p_\beta(\omega) := (p_\beta(a_1|\omega), ..., p_\beta(a_n|\omega))$.

By (5), the choice probabilities only depend on the state-wise utilities and even further only on state-wise utility differences with respect to some reference action. Using this, we show that the principal's problem is separable by state.

By a change of variables, the principal chooses directly the emotion-inclusive utility u instead of emotions m. The principal's maximization problem consists of an outer minimization problem regarding the implementation cost $\psi(u(a,\omega) - v(a,\omega))$ and an inner problem to maximize the payoff net of attention cost, which is subject to the exchangeability constraint:

$$\min_{u:A\times\Omega\to\mathbb{R}} \int \left(\sum_{\omega\in\Omega} \sum_{a\in A} \mu(\omega)p_\beta(a|\omega)\psi\big(u(a,\omega) - v(a,\omega)\big) \right) dF(\beta)$$

s.t.

$$u \in \operatorname*{arg\,max}_{u:A\times\Omega\to\mathbb{R}} \int \sum_{\omega\in\Omega} \mu(\omega) \left(\lambda\log(1/n) + \lambda H(p_\beta(\omega)) + \sum_{a\in A} p_\beta(a|\omega)v(a,\omega) \right) dF(\beta) \tag{6}$$

s.t. $\{a_1, ..., a_n\}$ is exchangeable with respect to u.

where $p_\beta(a|\omega)$ is given by (5) and $H(p_\beta(\omega)) = \sum_{a\in A} -p_\beta(a|\omega)\log(p_\beta(a|\omega))$ is the entropy of the conditional choice probabilities. We have used that under exchangeable payoffs, the unconditional

choice probabilities are all the same, and thus equal $1/n$. The constant entropy of the unconditional choice probabilities, $\lambda \log(1/n)$, can be dropped.

Inner Maximization Problem Turning to the inner maximization problem first, we can exchange the order of integration and summation in (6). We first relax the maximization problem by dropping the exchangeability constraint. Then, the inner objective is separable by state as $p_\beta(\omega)$ depends only on the vector of state-wise utilities, $u(\omega)$. Below, we show that any solution to the state-wise problem

$$\max_{u(\cdot,\omega):A\to\mathbb{R}} \int \left(\lambda H(p_\beta(\omega)) + \sum_{a\in A} p_\beta(a|\omega)v(a,\omega) \right) dF(\beta) \tag{7}$$

preserves the strict order of utilities and that a solution exists. Then, we go on to show that there is an exchangeable utility function that maximizes (7) state-by-state. By every state having positive probability, this implies that any solution to the constrained inner-maximization problem (including the exchangeability constraint) must maximize the state-wise problem in every state.

Lemma 2. *Any solution to the state-wise maximization problem (7) satisfies*

$$\forall a, a' \in A : v(a,\omega) > v(a',\omega) \Rightarrow u(a,\omega) > u(a',\omega).$$

Proof. If $v(a,\omega) > v(a',\omega)$ but $u(a,\omega) < u(a',\omega)$, then we can improve the principal's utility by exchanging $u(a,\omega)$ and $u(a',\omega)$, that is setting $u(a,\omega)$ to the original value of $u(a',\omega)$ and vice versa. For any type β, this would not change the entropy term while improving the expected payoff v. That is because for any type β, exchanging these two u's only exchanges the probabilities $p_\beta(a|\omega)$ and $p_\beta(a'|\omega)$ due to (5). The entropy $H(p_\beta(u(\omega)))$ remains unchanged for any type β because it is invariant under permutations of the probabilities, which is immediate from the definition of entropy. The expected payoff $\sum_{\hat{a}\in A} p_\beta(\hat{a}|\omega)v(\hat{a},\omega)$ would be improved because the conditional choice probabilities of $v(a,\omega)$ and $v(a',\omega)$ are exchanged, giving higher probability to the higher payoff. Hence, at a maximum, $v(a,\omega) > v(a',\omega)$ implies $u(a,\omega) \geq u(a',\omega)$.

Suppose for the sake of contraction that $v(a,\omega) > v(a',\omega)$ and $u(a,\omega) = u(a',\omega)$. Set $\tilde{u}(a,\omega) = u(a,\omega) + \varepsilon$ and $\tilde{u}(a',\omega) = u(a',\omega) - \varepsilon$ and $\tilde{u}(\hat{a},\omega) = u(\hat{a},\omega)$ for all $\hat{a} \neq a, a'$. Fix any type $\beta \in (0,1]$. At $\varepsilon = 0$, $p_\beta(a|u(\omega)) = p_\beta(a'|u(\omega))$ by (5). By symmetry of (5) in the u's, at $\varepsilon = 0$, for all $\hat{a} \neq a, a'$: $\frac{dp_\beta(\hat{a}|\omega)}{du(a,\omega)} = \frac{dp_\beta(\hat{a}|\omega)}{du(a',\omega)}$ and hence $\frac{dp_\beta(\hat{a}|\omega)}{d\varepsilon} = \frac{dp_\beta(\hat{a}|\omega)}{du(a,\omega)} - \frac{dp_\beta(\hat{a}|\omega)}{du(a',\omega)} = 0$ and $\frac{dp_\beta(a|\omega)}{d\varepsilon} = -\frac{dp_\beta(a'|\omega)}{d\varepsilon}$. By symmetry of entropy in the conditional choice probabilities, at $\varepsilon = 0$, $\frac{dH}{dp_\beta(a|\omega)} = \frac{dH}{dp_\beta(a'|\omega)}$ and hence $\frac{dH}{d\varepsilon} = \frac{dp_\beta(a'|\omega)}{d\varepsilon}\frac{dH}{dp_\beta(a'|\omega)} + \frac{dp_\beta(a|\omega)}{d\varepsilon}\frac{dH}{dp_\beta(a|\omega)} = 0$ by the equalities above. However, $v(a,\omega) - v(a',\omega) > 0$, so at $\varepsilon = 0$, the derivative of the expected payoff is positive by $\forall \hat{a} \neq a, a' : \frac{dp_\beta(\hat{a}|\omega)}{d\varepsilon} = 0, \frac{dp_\beta(a|\omega)}{d\varepsilon} > 0$ and $\frac{dp_\beta(a'|\omega)}{d\varepsilon} < 0$. Hence, increasing ε by some positive amount improves the principal's utility for all types β, which is why we must have $u(a,\omega) > u(a',\omega)$. $\square$

Because of the genericity Assumption 3, the Lemma implies that

$$v(a, \omega) > v(a', \omega) \Leftrightarrow u(a, \omega) > u(a', \omega).$$

Before we show existence, we insert the agent's state-dependent choice probabilities (5) to simplify the integrand of the state-wise objective:

$$-\lambda H(p_\beta(u(\omega))) + \sum_{a \in A} p_\beta(a|\omega) v(a, \omega)$$

$$= \sum_{a \in A} p_\beta(a|\omega) v(a, \omega) - \lambda p_\beta(a|\omega) \log(p_\beta(a|\omega))$$

$$= \sum_{a \in A} p_\beta(a|\omega) v(a, \omega) - \lambda \frac{e^{\frac{\beta}{\lambda} u(a, \omega)}}{\sum_{a' \in A} e^{\frac{\beta}{\lambda} u(a', \omega)}} \log \left(\frac{e^{\frac{\beta}{\lambda} u(a, \omega)}}{\sum_{a' \in A} e^{\frac{\beta}{\lambda} u(a', \omega)}} \right)$$

$$= \sum_{a \in A} p_\beta(a|\omega) \big(v(a, \omega) - \beta u(a, \omega) \big) + \lambda \log \left(\sum_{a' \in A} e^{\frac{\beta}{\lambda} u(a', \omega)} \right)$$

Thus, the state-wise objective, which we denote by $G(u(\omega))$, is

$$G(u(\omega)) = \int \left(\sum_{a \in A} p_\beta(a|u(\omega)) \big(v(a, \omega) - \beta u(a, \omega) \big) + \lambda \log \left(\sum_{a \in A} e^{\frac{\beta}{\lambda} u(a, \omega)} \right) \right) dF(\beta). \quad (8)$$

It will be useful to know the derivatives of G and $p_\beta(a|\omega)$ with respect to the components of u, which have simple expressions. Using

$$p_\beta(a|\omega) = \frac{d}{du(a, \omega)} \log \left(\sum_{a' \in A} e^{\frac{\beta}{\lambda} u(a', \omega)} \right),$$

which follows from (5), one can show by straightforward differentiation

$$\frac{dG(u(\omega))}{du(a, \omega)} = \int \left(\sum_{a' \in A} \frac{dp_\beta(a'|\omega)}{du(a, \omega)} \big(v(a', \omega) - \beta u(a', \omega) \big) \right) dF(\beta) \quad (9)$$

$$\frac{dp_\beta(a|\omega)}{du(a, \omega)} = \frac{\beta}{\lambda} p_\beta(a|\omega)(1 - p_\beta(a|\omega)) = \frac{\beta}{\lambda} p_\beta(a|\omega) \sum_{a' \neq a} p_\beta(a'|\omega) \quad (10)$$

$$\frac{dp_\beta(a'|\omega)}{du(a, \omega)} = -\frac{\beta}{\lambda} p_\beta(a'|\omega) p_\beta(a|\omega), \text{ for } a \neq a'. \quad (11)$$

Lemma 3. *The state-wise maximization problem (7) has a solution.*

Proof. We prove the statement by showing that any $u(\cdot, \omega) \in \mathbb{R}^A$ is weakly dominated (in the sense of achieving a lesser or equal objective) by some $\tilde{u}(\cdot, \omega) \in \mathbb{R}^A$ in a fixed compact subset of $\mathbb{R}^A$. Thus, we can restrict attention to this compact set, on which a maximum (which must be a global maximum) is obtained by continuity of the objective.

41

Without loss, assume $v(a_1, \omega) \leq ... \leq v(a_n, \omega)$. We can also assume without loss that $v(a_1, \omega) = u(a_1, \omega) = 0$. That is because from the state-wise objective (7) it is clear that shifting all $u(a, \omega)$ by the same constant does not affect choice probabilities and hence also not the objective. Similarly, shifting all $v(a, \omega)$ by some constant c changes the objective exactly by that constant since $\sum_{a \in A} p_\beta(a|\omega)(v(a, \omega) + c) = c + \sum_{a \in A} p_\beta(a|\omega)v(a, \omega)$. Hence, the constant does not affect the set of maximizers.

By Lemma 3, if for some a, a', $v(a, \omega) > v(a', \omega)$ and $u(a', \omega) > u(a, \omega)$, then we can exchange $u(a, \omega)$ and $u(a', \omega)$ and improve the principal's utility. If $v(a, \omega) = v(a', \omega)$ and $u(a', \omega) > u(a, \omega)$, then exchanging $u(a, \omega)$ and $u(a', \omega)$ does not change the principal's utility. Thus, every u that does not satisfy $u(a_1, \omega) \leq ... \leq u(a_n, \omega)$ is weakly dominated by a $\tilde{u}$ that does satisfy it. Thus, we can restrict attention to the set of utility-vectors u with $0 = u(a_1, \omega) \leq ... \leq u(a_n, \omega)$.

Finally, we bound $u(a_i, \omega)$ for $i = 2, ..., n$ from above by $u(a_{i-1}, \omega) + \frac{v(a_n, \omega)}{\inf \beta}$ ($\frac{1}{\inf \beta}$ exists by Assumption 1). Suppose for some $m \in \{2, ..., m\}$, $u(a_m, \omega)$ violated the bound. Then, for all $i \geq m, k \leq m - 1$:

$$u(a_i, \omega) \geq u(a_m, \omega) > u(a_{m-1}, \omega) + \frac{v(a_n, \omega)}{\inf \beta} \geq u(a_k, \omega) + \frac{v(a_n, \omega)}{\inf \beta}. \tag{12}$$

Using (9) through (11), the derivative of decreasing $u(a_m, \omega)$ through $u(a_n, \omega)$ by the same amount is

$$\left. \frac{dG(u(a_1, \omega), ..., u(a_{m-1}, \omega), u(a_m, \omega) - x, ..., u(a_n, \omega) - x)}{dx} \right|_{x=0} = -\sum_{i=m}^{n} \frac{dG(u)}{du(a_i, \omega)}$$

$$= \sum_{i=m}^{n} \int \frac{\beta}{\lambda} p_\beta(a_i|\omega) \left(\sum_{\substack{k=1, \\ k \neq i}}^{n} p_\beta(a_k|\omega) \Big[\big(v(a_k, \omega) - \beta u(a_k, \omega)\big) - \big(v(a_i, \omega) - \beta u(a_i, \omega)\big) \Big] \right) dF(\beta)$$

$$= \sum_{i=m}^{n} \int \frac{\beta}{\lambda} p_\beta(a_i|\omega) \left(\sum_{k=1}^{m-1} p_\beta(a_k|\omega) \Big[\big(v(a_k, \omega) - \beta u(a_k, \omega)\big) - \big(v(a_i, \omega) - \beta u(a_i, \omega)\big) \Big] \right) dF(\beta)$$

$$\geq \sum_{i=m}^{n} \int \frac{\beta}{\lambda} p_\beta(a_i|\omega) \left(\sum_{k=1}^{m-1} p_\beta(a_k|\omega) \Big[\beta\big(u(a_i, \omega) - u(a_k, \omega)\big) - v(a_n, \omega) \Big] \right) dF(\beta)$$

$$= \sum_{i=m}^{n} \int \frac{\beta}{\lambda} p_\beta(a_i|\omega) \beta \left(\sum_{k=1}^{m-1} p_\beta(a_k|\omega) \Big[u(a_i, \omega) - u(a_k, \omega) - \frac{v(a_n, \omega)}{\beta} \Big] \right) dF(\beta),$$

which is strictly positive since the term in the brackets is strictly positive by (12). Thus, we can decrease $u(a_m, \omega)$ through $u(a_n, \omega)$ until $u(a_i, \omega) \leq u(a_{i-1}, \omega) + \frac{v(a_n, \omega)}{\inf \beta}$, while improving the objective. Thus, each $u(a_i, \omega)$ can be bounded by $(i - 1)\frac{v(a_n, \omega)}{\inf \beta}$, which defines a compact set. $\square$

Lemma 4. *There exists an exchangeable utility function* $u : A \times \Omega \to \mathbb{R}$ *that maximizes (7) state-by-state.*

Proof. We construct a utility function that maximizes (7) state-by-state and which inherits ex-

changeability from v. The state-wise objective (7) is symmetric in the components of $v(\cdot, \omega)$. Thus, permuting the components of $v(\cdot, \omega)$ leads to a permutation of the components of the solution candidates $u(\cdot, \omega)$. Define an equivalence relation on Ω: two states are equivalent if their payoff-vectors $v(\cdot, \omega)$ are a permutation of each other. For each equivalence class, pick a representative state and a solution $u(\cdot, \omega)$ of the state-wise objective given that state's payoff-vector. Then, for other states ω' in the equivalence class, assign the analogous permutation of $u(\cdot, \omega)$, which is a solution given $v(\cdot, \omega')$. This procedure is well-defined because by Assumption 3, for any state ω, all components of $v(\cdot, \omega)$ are pairwise different. Thus, any permutation of $v(\cdot, \omega)$ can be obtained in a *unique* way by a sequence of pairwise permutations of components. The same holds for all components of any state-wise solution $u(\cdot, \omega)$, as by by Lemma 2 all components are pairwise different as well. On the equivalence class, any permutation of the state-wise solution of the representative state is thus as likely as any permutation by exchangeability of v. Doing the above procedure for every equivalence class, the resulting u is exchangeable. □

We have shown thus shown that the inner maximization problem separates into state-wise maximization of (7). Next, we show that the overall principal's problem separates by state.

State-Wise Principal's Problem Because the integral over states is a finite sum, we can exchange the order of integration in the implementation cost, which becomes

$$\sum_{\omega \in \Omega} \mu(\omega) \int \left(\sum_{a \in A} p_\beta(a|u(\omega)) \psi\big(u(a_i, \omega) - v(a_i, \omega)\big) \right) dF(\beta).$$

Because by constraint u is exchangeable, $p_\beta(a|u(\omega))$ depends only on the state-wise utility vector $u(\omega)$ by (5), so the outer objective separates by state. By the exchangeability assumption on v, one can make a symmetry argument that there is a state-wise minimizer u subject to it being among the state-wise maximizers of (7), such that the actions are exchangeable with respect to u. It follows that if a solution to the principal's problem exists, it must solve for every state $\omega \in \Omega$ the **state-wise principal's problem**:

$$\min_{u(\cdot, \omega):A \to \mathbb{R}} \int \left(\sum_{a \in A} p_\beta(a|u(\omega)) \psi\big(u(a, \omega) - v(a, \omega)\big) \right) dF(\beta)$$

s.t.

$$u(\cdot, \omega) \in \underset{u(\cdot, \omega):A \to \mathbb{R}}{\arg\max} \int \left(-\lambda H(p_\beta(u(\omega))) + \sum_{a \in A} p_\beta(a|u(\omega)) v(a, \omega) \right) dF(\beta)$$

Lemma 5. *A solution to the state-wise principal's problem exists. The solution is unique if the solution to the inner maximization problem is unique up to a common additive constant.*

Proof. We show again that we can restrict attention to a compact subset of $\mathbb{R}^A$, on which the minimum is obtained by continuity of the objective (ψ is convex implies that ψ is continuous).

First, we reparametrize $(u(a_1, \omega), ..., u(a_n, \omega))$ by

$$(u_1, \Delta u_2, ..., \Delta u_n) := (u(a_1, \omega), u(a_2, \omega) - u(a_1, \omega), ..., u(a_n, \omega) - u(a_1, \omega)).$$

Using $u(a_i, \omega) = u_1 + \Delta u_i$, the outer objective can be rewritten as

$$\sum_{i=1}^{n} \psi\big(u_1 + \Delta u_i - v(a_i, \omega)\big) \int p_\beta(a_i | u(\omega)) dF(\beta) \tag{13}$$

where $p_\beta(a_i | u(\omega))$ depends only on $(\Delta u_2, ..., \Delta u_n)$. By assumption, ψ is convex and minimized at zero. Thus, keeping $(\Delta u_2, ..., \Delta u_n)$ fixed, the objective is radially unbounded in u_1 as it is increasing for u_1 to the right of the interval

$$\Big[- \max_{i=2,...,n} \{\Delta u_i - v(a_i, \omega)\}, - \min_{i=2,...,n} \{\Delta u_i - v(a_i, \omega)\} \Big]$$

and decreasing when u_1 is to the left of the interval. Hence, because Δu_i is bounded by the proof in Lemma 3, this implies a bound on u_1. We know that u_1 does not affect the value of the inner objective because it does not affect the choice probabilities, so from the inner maximization problem, there is no other constraint on u_1. Further, by the proof of Lemma 3, $(\Delta u_2, ..., \Delta u_n)$ is in a compact set. By Berge's maximum theorem, the set of maximizers form a compact set themselves. Hence, we can restrict attention to a product of two compact sets, which is compact, on which the minimum is obtained.

For the second statement of the Lemma, if the solution to $(\Delta u_2, ..., \Delta u_n)$ from the inner maximization problem is unique, then the minimization problem above gives a unique solution for u_1 because equation (13) is a strictly convex function in u_1 keeping $(\Delta u_2, ..., \Delta u_n)$ fixed. $\qquad\square$

For $c \in \mathbb{R}$, let $\mathbf{c} := (c, ..., c) \in \mathbb{R}^A$ (with slight abuse of notation).

Lemma 6. $u(\cdot, \omega) \in \mathbb{R}^A$ *solves the state-wise principal's problem given the payoff vector* $v(\cdot, \omega) \in \mathbb{R}^A$ *if and only if* $u(\cdot, \omega) + \mathbf{c}$ *does so given payoff vector* $v(\cdot, \omega) + \mathbf{c}$.

Proof. As noted in Lemma 3, shifting all components of $v(\cdot, \omega)$ or $u(\cdot, \omega)$ by constants does not affect the inner maximization problem. The solutions to the inner maximizaton problem given $v(\cdot, \omega)$ and $v(\cdot, \omega) + \mathbf{c}$ are the same as this change only moves the principal's utility by c independent of the agent's choice. And, if $u(\cdot, \omega)$ is a solution to the inner maximization problem, then $u(\cdot, \omega) + \mathbf{c}$ is, too, since shifting all components of $u(\cdot, \omega)$ by a constant does not affect the principal's utility as it does not affect the agent's choice probabilities. Thus, if $u(\cdot, \omega)$ is in the constraint set defined by the inner maximization problem given $v(\cdot, \omega)$, then $u(\cdot, \omega) + \mathbf{c}$ is in the said set given $v(\cdot, \omega) + \mathbf{c}$.

Second, the outer objective is unchanged if $\mathbf{c}$ is added to both $u(\cdot, \omega)$ and $v(\cdot, \omega)$ because the agent's choice probabilities are unaffected and $\psi(u(a_i, \omega) - v(a_i, \omega))$ is unaffected. If $u(\cdot, \omega)$ minimizes the implementation cost given $v(\cdot, \omega)$, then $u(\cdot, \omega) + \mathbf{c}$ minimizes it given $v(\cdot, \omega) + \mathbf{c}$ over the original constraint set plus $\mathbf{c}$. However, this is the actual constraint set, by the above. The other direction follows from shifting by $-\mathbf{c}$. $\qquad\square$

Solution to Principal's Problem

Lemma 7. *There exists a solution to the principal's problem such that there is a function R, such that for all $(a, \omega) \in A \times \Omega$,*

$$u(a, \omega) = v(a, \omega) + R(\{v(a, \omega) - v(a', \omega)\}_{a' \in A}).$$

Proof. Define a partition $\mathbb{P}$ on the states based on the equivalence relation that the payoff vector $v(\cdot, \omega) \in \mathbb{R}^A$ of two states are the same up to permutations of its components, shifting all components by a constant, and combinations of both. Formally, $\omega \sim \omega'$ if there exists a permutation $\tilde{v}(\cdot, \omega)$ of $v(\cdot, \omega)$ and a $c \in \mathbb{R}$ such that $v(\cdot, \omega) = \tilde{v}(\cdot, \omega) + \mathbf{c}$. We proceed by constructing $u(\cdot, \omega)$ on each partition element.

Let $P \in \mathbb{P}$ be some partition element and ω be some state in P. Define $u(\cdot, \omega)$ to be a solution to the state-wise principal's problem, which exists by Lemma 5. So, for each other state ω' in P, there exist a $c \in \mathbb{R}$ and a permutation $\tilde{v}(\cdot, \omega)$ of the payoff vector $v(\cdot, \omega)$, such that $v(\cdot, \omega') = \mathbf{c} + \tilde{v}$. Moreover, the constant c and the way $v(\cdot, \omega)$ is permuted are unique, the latter because by Assumption 3. Define $u(\cdot, \omega')$ to be the analogous permutation of $u(\cdot, \omega)$ plus $\mathbf{c}$. This gives a solution to the state-wise principal's problem under ω' because the state-wise principal's problem is the same up to this permutation and by Lemma 6. Because the actions are exchangeable with respect to $v(\cdot, \omega)$, each state in partition element in P that has $c = 0$ relative to ω (that is a "pure permutation") has the same probability. Hence, by construction $u(\cdot, \omega)$ is a pure permutation on these elements, so the actions are exchangeable with respect to $u(\cdot, \omega)$. Thus, the constructed $u(\cdot, \omega)$ solves the principal's problem.

Further, $u(\cdot, \omega)$ is a function only of the payoff vector $v(\cdot, \omega)$. Denote this function by $M : \mathbb{R}^A \to \mathbb{R}^n$ and by M_i the $i-th$ component of M. Define for all $(v_1, ..., v_{n-1}) \in \mathbb{R}^A$

$$\tilde{R}(v_1, ..., v_{n-1}) := M_n(v_1, ..., v_{n-1}, 0).$$

By construction, a permutation of the arguments of M permutes the image in an analogous way. Formally, for all pairwise different $i, j, k \in \{1, ..., n\}$

$$M_i(v_1, ..., v_i, ..., v_j, ..., v_n) = M_j(v_1, ..., v_j, ..., v_i, ..., v_n)$$
$$M_k(v_1, ..., v_i, ..., v_j, ..., v_n) = M_k(v_1, ..., v_j, ..., v_i, ..., v_n)$$

This implies that $\tilde{R}$ is invariant under relabeling of its components, so we can define

$$R(\{v_1, ..., v_n\}) := \tilde{R}(-v_1, ..., -v_n).$$

We have by Lemma 6,

$$u(a_i, \omega) = M_i(v(a_1, \omega), ..., v(a_2, \omega))$$

$$= v(a_i, \omega) + M_i(v(a_1, \omega) - v(a_i, \omega), ..., v(a_i, \omega) - v(a_i, \omega), ..., v(a_2, \omega) - v(a_i, \omega))$$

$$= v(a_i, \omega) + M_n(v(a_1, \omega) - v(a_i, \omega), ..., v(a_2, \omega) - v(a_i, \omega), 0)$$

$$= v(a_i, \omega) + \tilde{R}(v(a_1, \omega) - v(a_i, \omega), ..., v(a_2, \omega) - v(a_i, \omega))$$

$$= v(a_i, \omega) + R(\{v(a_i, \omega) - v(a', \omega)\}_{a' \in A \setminus \{a_i\}})$$

□

Note that the constructed solution to the state-wise principal's problem depends only on the state-wise actual and counterfactual payoffs as well as the distribution over β and λ. Thus, the solution does not depend on the whole on the decision-problem that the principal delegates.

D   Proof of Theorem 1

Proof. By Theorem 2 there is a solution with

$$u(a, \omega) = v(a, \omega) + R(v(a, \omega) - v(a_-, \omega)), \tag{14}$$

if $v(a, \omega) - v(a_-, \omega) \neq 0$ for all states (because we have assumed strict preferences of the principal for that theorem). First, we show that this statement extends to cases where $v(a, \omega) - v(a_-, \omega) = 0$ for some states. Next, we show that any two such solutions have R and Q that are unique up to a measure zero set in $\mathbb{R}$ (with respect to the Lebesgue measure). Then, we show that any such solution satisfies Properties 1 and 2.

Extending Theorem 1   By Theorem 2, a solution to the principal's problem exists and every solution must, in every state, solve the state-wise principal's problem, which is

$$\min_{u(\cdot, \omega): A \to \mathbb{R}} \int \left(\sum_{i=1}^{2} p_\beta(a_i | u(\omega)) \psi\big(u(a_i, \omega) - v(a_i, \omega)\big) \right) dF(\beta)$$

s.t.

$$u(\cdot, \omega) \in \arg\max_{u(\cdot, \omega): A \to \mathbb{R}} \int \left(\lambda H(p_\beta(u(\omega))) + \sum_{i=1}^{2} p_\beta(a_i | u(\omega)) v(a_i, \omega) \right) dF(\beta)$$

If $v(a_1, \omega) = v(a_2, \omega)$, the payoff $\sum_{i=1}^{2} p_\beta(a_i | u(\omega)) v(a_i, \omega)$ is unaffected by the conditional choice probabilities $p_\beta(a_i | u(\omega))$. The entropy term is maximized if

$$p_\beta(a_1 | u(\omega)) = p_\beta(a_2 | u(\omega)) = 1/2,$$

which is achieved for all types if and only if $u(a_1, \omega) = u(a_2, \omega)$. In that case, because the implementation cost is minimized at zero, the unique solution is

$$u(a_1, \omega) = u(a_2, \omega) = v(a_1, \omega) = v(a_2, \omega).$$

Take our solution from Theorem 2 for states where $v(a_1, \omega) \neq v(a_2, \omega)$, we can extend it for states where $v(a_1, \omega) = v(a_2, \omega)$ and maintain exchangeability, so it must be a solution.

Uniqueness By (8), the inner objective can be rewritten as

$$\int \left(\sum_{i=1}^{2} p_\beta(a_i | u(\omega)) \big(v(a_i, \omega) - \beta u(a_i, \omega) \big) + \lambda \log \left(\sum_{i=1}^{2} e^{\frac{\beta}{\lambda} u(a_i, \omega)} \right) \right) dF(\beta).$$

Using $p_\beta(a_1 | \omega) = 1 - p_\beta(a_2 | \omega)$, we can rewrite the objective as

$$\int p_\beta(a_2 | u(\omega)) \Big[v(a_2, \omega) - v(a_1, \omega) - \beta \big(u(a_2, \omega) - u(a_1, \omega) \big) \Big] + v(a_1, \omega) - \beta u(a_1, \omega)$$

$$+ \lambda \left(\frac{\beta}{\lambda} u(a_1, \omega) + \log \left(1 + e^{\frac{\beta}{\lambda}(u(a_2, \omega) - u(a_1, \omega))} \right) \right) dF(\beta)$$

Defining $\Delta u(\omega) = u(a_2, \omega) - u(a_1, \omega)$ and $\Delta v(\omega) = v(a_2, \omega) - v(a_1, \omega)$, and subtracting the constant $v(a_1, \omega)$, we obtain that the objective is

$$\int \left(p_\beta(a_2 | u(\omega)) \big(\Delta v(\omega) - \beta \Delta u(\omega) \big) + \lambda \log \left(1 + e^{\frac{\beta}{\lambda} \Delta u(\omega)} \right) \right) dF(\beta) \tag{15}$$

Note that $p_\beta(a_2 | u(\omega))$ depends only on $\Delta u(\omega)$ by (5), so with slight abuse of notation, we write $p_\beta(a_2 | \Delta u(\omega))$. The first-order condition is

$$\int \left(\frac{dp_\beta(a_2 | \Delta u(\omega))}{d\Delta u(\omega)} \big(\Delta v(\omega) - \beta \Delta u(\omega) \big) \right) dF(\beta) = 0, \tag{16}$$

which can be rewritten as

$$\Delta u(\omega) = \Delta v(\omega) \left(\frac{\int \frac{dp_\beta(a_2 | \Delta u(\omega))}{d\Delta u(\omega)} \beta dF(\beta)}{\int \frac{dp_\beta(a_2 | \Delta u(\omega))}{d\Delta u(\omega)} dF(\beta)} \right)^{-1}. \tag{17}$$

as used in the main text.

The solution $\Delta u(\omega)$ need not be unique. However, using the supermodularity of the objective in $\Delta u(\omega)$ and $\Delta v(\omega)$, we show that there is a unique maximizer $\Delta u(\omega)$ for almost all $\Delta v(\omega)$.

Lemma 8. *The set of solutions $\Delta u(\omega)$ as a function of $\Delta v(\omega)$ are increasing in the strong set order. For almost all $\Delta v(\omega)$, the solution $\Delta u(\omega)$ is unique.*

Proof. The objective (15) has a positive cross-derivative in $\Delta v(\omega)$ and $\Delta u(\omega)$, which, using (16), is

$$\int \frac{dp_\beta(a_2|\Delta u(\omega))}{d\Delta u(\omega)} dF(\beta) > 0.$$

It is positive because $p_\beta(a_2|\Delta u(\omega))$ is strictly increasing in $\Delta u(\omega)$ for all $\beta > 0$. Thus, by Topkis characterization theorem (Topkis, 1978), the objective is supermodular in $\Delta u(\omega)$ and $\Delta v(\omega)$. By Topkis theorem (ibd.), the set of $\Delta u(\omega)$ that maximize the objective as a function of $\Delta v(\omega)$ is increasing in the strong set order. Since the cross-partial is strictly positive, an even stronger statement holds: if $\Delta u'' > \Delta u'$ are both maximizers at $\Delta v(\omega)$, then at any $\Delta v(\omega)' > \Delta v(\omega)$, $\Delta u''$ achieves a strictly higher payoff than $\Delta u'$. That is because the derivative of the objective between $\Delta u'$ and $\Delta u''$ is strictly higher at any point, so the integral over the derivative is now positive, where it was zero before. Thus, any selection function of maximizers as a function of $\Delta v(\omega)$ must have a jump discontinuity wherever the set of maximizers is not a singleton. Since monotonic functions can have at most countably many jump discontinuities, there are at most countably many $\Delta v(\omega)$ for which there is not a unique maximizer. Thus, for almost all $\Delta v(\omega)$, the solution is unique. $\square$

By Lemma 5 and 8, the solution to the outer minimization problem subject to the inner maximization problem is unique for almost all $\Delta v(\omega)$. Thus, $R(x)$ is unique for almost all $x \in \mathbb{R}$ and thereby also $Q(x) := x + R(x) - R(-x)$, which equals $\Delta u(\omega)$ given $\Delta v(\omega) = x$.

Property 1 Above, we have already argued that $R(0) = 0$. First, we show that $R(x) \geq 0$ for $x > 0$. Note that equation (17) show that $Q(x) - x > 0$ when $x > 0$. Thus, if $R(x)$ was negative then $R(x)$ and $R(-x)$ were negative by $R(x) - R(-x) = Q(x) - x$. Then, we could increase $R(x)$ and $R(-x)$ by the same small amount (thus maintaining $Q(x)$ being a solution to the state-wise maximization problem) and decrease the implementation cost by ψ being strictly convex and minimal at zero. Analogously, one can show that $R(x) \leq 0$ if $x < 0$.

If $R(x)$ was zero, then and $R(-x)$ would be negative by $R(x) - R(-x) = Q(x) - x$. Because ψ is differentiable and has a global minimum at zero, it must have derivative 0 at 0. Then, we could increase $R(x)$ and $R(-x)$ by the some small enough amount (thus maintaining $Q(x)$ being a solution to the state-wise maximization problem) and decrease the implementation cost, by ψ being strictly convex and having derivative 0 at 0. Analogously, one can show that $R(x) < 0$ if $x < 0$.

Property 2 As argued in the main text, Property 2 is equivalent to $Q(x)/x$ being strictly increasing on $\mathbb{R}_{>0}$. We show that any $Q(x) := x + R(x) - R(-x)$ satisfies this property. Rearranging (17), we get

$$\frac{Q(x)}{x} = \left(\frac{\int \frac{dp_\beta(a_2|Q(x))}{dQ(x)} \beta dF(\beta)}{\int \frac{dp_\beta(a_2|Q(x))}{dQ(x)} dF(\beta)} \right)^{-1}. \tag{18}$$

48

Because β has a discrete distribution by Assumption 1, the integral $\int g(\beta) dF(\beta)$ over some measurable function $g(\beta)$ equals $\sum_{\beta \in \text{supp}(\beta)} g(\beta) f(\beta)$ with probability mass function f. The term in parentheses in (18) is the expectation of β with respect to the probability mass function

$$\frac{\frac{dp_\beta(a_2|Q(x))}{dQ(x)} f(\beta)}{\int \frac{dp_b(a_2|Q(x))}{dQ(x)} dF(b)}.$$

Since we are considering $x > 0$, we have $Q(x) > 0$. Lemma 9 below shows that this probability mass function undergoes a strict likelihood-ratio shift downwards as $Q(x)$ increases for $Q(x) > 0$. Because $Q(x)$ is strictly increasing by Lemma 8 and because strict likelihood-ratio dominance implies strict first-order stochastic dominance, this implies that the term in parentheses in (18) is strictly decreasing in x. Thus, $Q(x)/x$ is strictly increasing in x.

Lemma 9. *The probability mass function*

$$\frac{\frac{dp_\beta(a_2|Q(x))}{dQ(x)} f(\beta)}{\int \frac{dp_b(a_2|Q(x))}{dQ(x)} dF(b)}$$

undergoes a strict likelihood-ratio shift downwards as $Q(x)$ increases for $Q(x) > 0$.

Proof. We need to show that for all $1 \geq \alpha > \beta > 0$,

$$\frac{\left(\frac{\frac{dp_\alpha(a_2|Q(x))}{dQ(x)} f(\alpha)}{\int \frac{dp_b(a_2|Q(x))}{dQ(x)} dF(b)} \right)}{\left(\frac{\frac{dp_\beta(a_2|Q(x))}{dQ(x)} f(\beta)}{\int \frac{dp_b(a_2|Q(x))}{dQ(x)} dF(b)} \right)}$$

is decreasing in $Q(x)$. Canceling out the integration constants in the denominators, this term is equal to

$$\frac{\frac{dp_\alpha(a_2|Q(x))}{dQ(x)} f(\alpha)}{\frac{dp_\beta(a_2|Q(x))}{dQ(x)} f(\beta)} = \frac{\left(\frac{\frac{\alpha}{\lambda} e^{\frac{\alpha}{\lambda} Q(x)} f(\alpha)}{\left(1+e^{\frac{\alpha}{\lambda} Q(x)}\right)^2} \right)}{\left(\frac{\frac{\beta}{\lambda} e^{\frac{\beta}{\lambda} Q(x)} f(\beta)}{\left(1+e^{\frac{\beta}{\lambda} Q(x)}\right)^2} \right)} = \frac{\alpha f(\alpha)}{\beta f(\beta)} e^{\frac{\alpha-\beta}{\lambda} Q(x)} \left(\frac{1+e^{\frac{\beta}{\lambda} Q(x)}}{1+e^{\frac{\alpha}{\lambda} Q(x)}} \right)^2$$

Ignoring the constant $\frac{\beta f(\beta)}{\alpha f(\alpha)}$, the term is decreasing in $Q(x)$ if the logarithm of the term is decreasing in $Q(x)$.

$$\log\left(e^{\frac{\alpha-\beta}{\lambda} Q(x)} \left(\frac{1+e^{\frac{\beta}{\lambda} Q(x)}}{1+e^{\frac{\alpha}{\lambda} Q(x)}} \right)^2 \right) = (\alpha-\beta)\frac{Q(x)}{\lambda} + 2\log\frac{1+e^{\frac{\beta}{\lambda} Q(x)}}{1+e^{\frac{\alpha}{\lambda} Q(x)}}$$

Differentiating with respect to $Q(x)$:

$$\frac{d}{dQ(x)}\left((\alpha-\beta)\frac{Q(x)}{\lambda}+2\log\frac{1+e^{\frac{\beta}{\lambda}Q(x)}}{1+e^{\alpha Q(x)/\lambda}}\right)=\frac{\alpha-\beta}{\lambda}+2\frac{\frac{d}{dQ(x)}\frac{1+e^{\frac{\beta}{\lambda}Q(x)}}{1+e^{\frac{\alpha}{\lambda}Q(x)}}}{\frac{1+e^{\frac{\beta}{\lambda}Q(x)}}{1+e^{\frac{\alpha}{\lambda}Q(x)}}}$$

$$=\frac{\alpha-\beta}{\lambda}+2\frac{\frac{\beta}{\lambda}e^{\frac{\beta}{\lambda}Q(x)}(1+e^{\frac{\alpha}{\lambda}Q(x)})-\frac{\alpha}{\lambda}e^{\frac{\alpha}{\lambda}Q(x)}(1+e^{\frac{\beta}{\lambda}Q(x)})}{\left(1+e^{\frac{\alpha}{\lambda}Q(x)}\right)^2}\frac{1+e^{\frac{\alpha}{\lambda}Q(x)}}{1+e^{\frac{\beta}{\lambda}Q(x)}}$$

$$=\frac{\alpha-\beta}{\lambda}+2\frac{\beta-\alpha}{\lambda}\frac{e^{\frac{\alpha+\beta}{\lambda}Q(x)}}{(1+e^{\frac{\beta}{\lambda}Q(x)})(1+e^{\frac{\alpha}{\lambda}Q(x)})}+2\frac{\frac{\beta}{\lambda}e^{\frac{\beta}{\lambda}Q(x)}-\frac{\alpha}{\lambda}e^{\frac{\alpha}{\lambda}Q(x)}}{(1+e^{\frac{\alpha}{\lambda}Q(x)})(1+e^{\frac{\beta}{\lambda}Q(x)})}$$

Multiply by $\lambda(1+e^{\frac{\beta}{\lambda}Q(x)})(1+e^{\frac{\alpha}{\lambda}Q(x)})>0$ gives

$$(\alpha-\beta)(1+e^{\frac{\beta}{\lambda}Q(x)})(1+e^{\frac{\alpha}{\lambda}Q(x)})+2(\beta-\alpha)e^{\frac{\alpha+\beta}{\lambda}Q(x)}+2(\beta e^{\frac{\beta}{\lambda}Q(x)}-\alpha e^{\frac{\alpha}{\lambda}Q(x)})$$

$$=(\alpha-\beta)+(\alpha+\beta)e^{\frac{\beta}{\lambda}Q(x)}-(\alpha+\beta)e^{\frac{\alpha}{\lambda}Q(x)}+(\beta-\alpha)e^{\frac{\alpha+\beta}{\lambda}Q(x)}$$

$$=(\alpha+\beta)(e^{\frac{\beta}{\lambda}Q(x)}-e^{\frac{\alpha}{\lambda}Q(x)})+(\alpha-\beta)(1-e^{\frac{\beta}{\lambda}Q(x)}e^{\frac{\alpha}{\lambda}Q(x)}).$$

By $\alpha,\beta>0$ and $Q(x)>0$, we have $e^{\frac{\alpha}{\lambda}Q(x)}>e^{\frac{\beta}{\lambda}Q(x)}$ and hence the first term is negative. By $\alpha>\beta$ and $Q(x)>0$, we have $e^{\frac{\beta}{\lambda}Q(x)}>1$, $e^{\frac{\alpha}{\lambda}Q(x)}>1$, and $e^{\frac{\beta}{\lambda}Q(x)}e^{\frac{\alpha}{\lambda}Q(x)}>1$. Hence, the second term is negative. This establishes the negative derivative in $Q(x)$. $\square$

Thus, Q and R of the regret solution satisfy

$$Q(x)\in\underset{y\in\mathbb{R}}{\arg\max}\int\frac{e^{\beta y/\lambda}}{1+e^{\beta y/\lambda}}(x-\beta y)+\lambda\log(1+e^{\beta y/\lambda})dF(\beta)$$

$$R(x)=\underset{R\in\mathbb{R}}{\arg\min}\left(\int p_\beta(Q(x))dF(\beta)\right)\psi(R)+\left(1-\int p_\beta(Q(x))dF(\beta)\right)\psi(x+R-Q(x))$$

where $p_\beta(Q(x)):=\frac{e^{\frac{\beta}{\lambda}Q(x)}}{1+e^{\frac{\beta}{\lambda}Q(x)}}$. $\square$

E Proof of Proposition 1

Our Property 2 is slightly weaker than the standard convexity assumption of Q, or equivalently the assumption that R is decreasingly concave (Bell, 1982; Loomes and Sugden, 1982). One can see how they relate using the terminology of production functions. Viewing Q as a production function for the sake of the analogy, convexity corresponds to increasing *marginal* product, while Property 2 corresponds to increasing *average* product. Under differentiability, increasing *average* product is equivalent to the marginal product being greater than the average product, in contrast to convexity which requires that the marginal product at x is greater than *any* marginal product up to x. Loomes and Sugden (1982) motivate the *convex* regret-rejoicing function Q through its implication of the Allais paradox, the common-ratio effect, the isolation effect, and under linear

50

utility over money, the reflection effect, and simultaneous gambling and insurance. They show that the reflection effect holds under regret theory independent of assumptions on Q and that the isolation effect is equivalent to the common-ratio effect. Thus, we focus on the (a) Allais paradox, (b) the common-ratio effect, and (c) simultaneous gambling and insurance below.[34] Proposition 1 shows that Property 2 is in fact equivalent to the conjunction of (a) to (c) in the context of regret theory.

Note that regret theory was developed for decisions without taking into account attention. In particular, it has been applied to decisions between lotteries, where the nature of uncertainty is such that the state cannot be learned by paying attention. One way to reconcile this with our model is to assume that the agent's internal self-control mechanism treats those decisions *as if* the state could be learned and triggers regret and rejoicing regardless.[35]

The above-mentioned behavioral effects are formulated in terms of lotteries. To identify lotteries with actions in our present framework, let $x(a, \omega)$ denote the monetary outcome of action a under state ω. Given the prior μ over states, an action induces a lottery over monetary amounts $x(a, \omega)$. Following Loomes and Sugden (1982), lotteries can be identified with any action that induces such lottery given the belief μ. Two independent lotteries can be identified with two actions such that the two distributions over monetary amounts induced by the prior are independent. Let $(x_1, p; x_2, 1-p)$ denote a lottery that gives monetary outcome x_1 with probability p and monetary outcome x_2 with probability $1 - p$. Further, let (x, p) be shorthand for $(x, p; 0, 1-p)$.

Like Loomes and Sugden (1982), we consider choice sets of two actions, $A = \{a_1, a_2\}$. Let $v : \mathbb{R} \to \mathbb{R}$ be the decision-maker's increasing utility function over money. According to regret theory, the decision-maker's utility of action $a_i \in A$ is given by

$$\sum_{\omega \in \Omega} v(x(a_i, \omega)) + R\big(v(x(a_i, \omega)) - v(x(a_{-i}, \omega))\big)$$

for some regret-rejoicing function $R : \mathbb{R} \to \mathbb{R}$. Equivalently, let $Q(x) = x + R(x) - R(-x)$ and preferences between actions be given by

$$a_1 \succsim a_2 \Leftrightarrow \sum_{\omega \in \Omega} \mu(\omega) Q\big(v(x(a_1, \omega)) - v(x(a_2, \omega))\big) \geq 0 \tag{19}$$

where $Q(x) = x + R(x) - R(-x)$. At this point, we impose no assumptions on R or Q.

The following definitions are taken from Loomes and Sugden (1982).

Definition 6 (Allais Paradox). *Let $X_1 = (x_1, p_1; x_2, \alpha)$ and $X_2 = (x_2, p_2 + \alpha)$ be independent lotteries where $1 \geq p_2 > p_1 > 0$ and $(1 - p_2) \geq \alpha \geq 0$. If there exists some probability $\overline{\alpha}$ such that*

---

[34]Bell (1982) observes that Property 2 is sufficient for simultaneous gambling and insurance and for rejection of probabilistic insurance. Loomes and Sugden (1987) demonstrate that superadditivity of regret, which we show below follows from Property 2, predicts preferences reversals.

[35]Regret and rejoicing can be understood as rational *rules*, that perform well in most environments, in which arguably some learning about the state is possible. See for example Heiner (1983), Aumann (1997), and Halevy and Feltkamp (2005) for related discussions on rule rationality.

$X_1 \sim X_2$ when $\alpha = \overline{\alpha}$, then (i) if $x_1 > x_2 > 0$, then $\alpha < \overline{\alpha} \Rightarrow X_1 \succ X_2$ and $\alpha > \overline{\alpha} \Rightarrow X_1 \prec X_2$ and (ii) if $0 > x_2 > x_1$, then $\alpha < \overline{\alpha} \Rightarrow X_1 \prec X_2$ and $\alpha > \overline{\alpha} \Rightarrow X_1 \succ X_2$.

Definition 7 (Common Ratio Effect). *Let $X_1 = (x_1, \lambda p)$ and $X_2 = (x_2, p)$ be independent lotteries, where $1 \geq p > 0$ and $1 > \lambda > 0$. If there exists some probability $\overline{p}$ such that $X_1 \sim X_2$ when $p = \overline{p}$, then (i) if $x_1 > x_2 > 0$, then $p < \overline{p} \Rightarrow X_1 \succ X_2$ and $p > \overline{p} \Rightarrow X_1 \prec X_2$ and (ii) if $0 > x_2 > x_1$, then $p < \overline{p} \Rightarrow X_1 \prec X_2$ and $p > \overline{p} \Rightarrow X_1 \succ X_2$.*

Definition 8 (Simultaneous Gambling and Insurance). *Let v be linear in money. Let $X_1 = (0, 1)$ and $X_2(x, p; -px/(1-p), 1-p)$ be independent lotteries where $0 < p < 1$ and $x > 0$. (i) For $p > 0.5$, $X_1 \succ X_2$, (ii) for $p < 0.5$, $X_1 \prec X_2$, and for (iii) for $p = 0.5$, $X_1 \sim X_2$.*

Proposition 1. *Suppose the decision-maker has preferences between lotteries given by (19). The decision-maker's preferences exhibit the Allais paradox, common-ratio effect, and simultaneous gambling and insurance if and only if $Q(x)/x$ is strictly increasing on $\mathbb{R}_{>0}$.*

Proof. Loomes and Sugden (1982) show that the Allais paradox and the common-ratio effect are equivalent to strict superadditivity,

$$\forall x_1, x_2 : 0 < x_2 < x_1 \rightarrow Q(x_1) - Q(x_1 - x_2) - Q(x_2) > 0, \tag{20}$$

and that simultaneous gambling and insurance is equivalent to

$$\forall x > 0 : Q(\frac{p}{1-p}x) \gtreqless \frac{p}{1-p}Q(x) \quad \text{if } p \gtreqless 0.5. \tag{21}$$

We show that (20) and (21) are implied by $Q(x)/x$ being strictly increasing on $\mathbb{R}_{>0}$ and that (21) implies that $Q(x)/x$ is strictly on $\mathbb{R}_{>0}$.

To show that (20) is implied by Property 2, let $x_1 > x_2 > 0$ and define $Q_1 := Q(x_1)$, $Q_2 := Q(x_2)$, and $Q_{12} := Q(x_1 - x_2)$. We can assume without loss $0 < x_1 - x_2 < x_2 < x_1$. That is because if $0 < x_2 < x_1 - x_2 < x_1$, we can simply exchange x_2 and $x_1 - x_2$ in (20). By $Q(x)/x$ being strictly increasing on $\mathbb{R}_{>0}$:

$$\frac{Q_1}{x_1} > \frac{Q_2}{x_2}$$
$$\Rightarrow Q_1 x_2 > Q_2 x_1$$
$$\Rightarrow Q_1 x_2 - Q_2 x_2 > Q_2 x_1 - Q_2 x_2$$
$$\Rightarrow (Q_1 - Q_2) x_2 > Q_2 (x_1 - x_2)$$
$$\Rightarrow \frac{Q_1 - Q_2}{x_1 - x_2} > \frac{Q_2}{x_2}$$

Using again that $Q(x)/x$ is strictly increasing on $\mathbb{R}_{>0}$,

$$\frac{Q_2}{x_2} > \frac{Q_{12}}{x_1 - x_2}.$$

Together,

$$\frac{Q_1 - Q_2}{x_1 - x_2} > \frac{Q_2}{x_2} > \frac{Q_{12}}{x_1 - x_2} \Rightarrow Q_1 - Q_{12} - Q_2 > 0.$$

Equation (21) follows from $Q(x)/x$ strictly increasing on $\mathbb{R}_{>0}$ by

$$p > 0.5 \Rightarrow \frac{p}{1-p} > 1 \Rightarrow \frac{Q(\frac{p}{1-p}x)}{\frac{p}{1-p}x} > \frac{Q(x)}{x} \Rightarrow Q(\frac{p}{1-p}x) > \frac{p}{1-p}Q(x),$$

and analogously for $p < 0.5$. When $p = 0.5$, $\frac{p}{1-p} = 1$ and the statement trivially holds.

For the other direction, suppose $Q(x)/x$ was not strictly increasing on $\mathbb{R}_{>0}$. Then, there exists $0 < x < y$ such that $\frac{Q(y)}{y} \leq \frac{Q(x)}{x}$. Let $p = \frac{y/x}{1+y/x} \in (1/2, 1)$. Then, $\frac{p}{1-p} = \frac{y}{x}$ and thus

$$\frac{Q(\frac{p}{1-p}x)}{\frac{p}{1-p}x} \leq \frac{Q(x)}{x} \Rightarrow Q(\frac{p}{1-p}x) \leq \frac{p}{1-p}Q(x)$$

despite $p > 1/2$ violating (21). $\qquad \square$

F Proof of Corollary 1

Proof. Regarding the action-implementation stage, we need to establish that given the information the agent of type β acquires, he takes the principal's preferred action. That is, we show that the posterior of type β under some action $a_i \in A$ makes a_i the optimal choice for the principal. This posterior is proportional to the product of the prior and the likelihood, the latter of which is given by the state-dependent choice probabilities. Consider the posterior probability of some emotion-*exclusive* payoff vector $(v(a_i, \omega), v(a-i, \omega)) = (v_1, v_2)$ and its permutation (v_2, v_1) with $v_1 > v_2$. Both payoff vectors have the same prior probability by exchangeability. Because emotions preserve the state-wise ordering of payoffs, the conditional choice probability of any agent β of action a_i is higher than the one of a_{-i} by (5). Thus, the posterior probability of (v_1, v_2) conditional on a_i is higher than conditional on a_{-i}, contributing to a higher expected utility for the principal of a_i than a_{-i}. Integrating over all pairs of payoff vectors and their permutation, the expected utility of action a_i with respect to v is higher than the one of action a_{-i}. $\qquad \square$

G Proof of Corollary 2

Proof. Let $x > 0$ and denote $p_\beta(Q(x)) := \frac{e^{\frac{\beta}{\lambda}Q(x)}}{1+e^{\frac{\beta}{\lambda}Q(x)}}$. From the proof of Theorem 1, regret and rejoicing, $R(-x)$ and $R(x)$, minimize the implementation cost subject to being consistent with $Q(x)$:

$$\min_{R(x),R(-x)\in\mathbb{R}} \left(\int p_\beta(Q(x))dF(\beta)\right)\psi(R(x)) + \left(1 - \int p_\beta(Q(x))dF(\beta)\right)\psi(R(-x))$$

$$\text{s.t.}$$

$$x + R(x) - R(-x) = Q(x)$$

53

That x is positive implies $Q(x) > 0$, which in turn implies that for all $\beta > 0$:

$$p_\beta(Q(x)) = \frac{e^{\frac{\beta}{\lambda}Q(x)}}{1 + e^{\frac{\beta}{\lambda}Q(x)}} > \frac{1}{2} > 1 - p_\beta(Q(x))$$

$$\Rightarrow \int p_\beta(Q(x)) dF(\beta) > 1 - \int p_\beta(Q(x)) dF(\beta)$$

Thus, in the implementation-cost minimization problem, the weight on $\psi(R(x))$ is greater than the weight on $\psi(R(-x))$. Suppose for the sake of contradiction that we had $R(x) > |R(-x)|$. By convexity of ψ and symmetry around 0, there is an $\varepsilon > 0$ such that the implementation cost decreases if $R(x)$ and $R(-x)$ are both lowered by ε, which maintains the constraint $x + R(x) - R(-x) = Q(x)$. $\qquad\square$

H Proof of Corollary 3

Proof. We show, first, that decreasing λ increases $Q(x)$ for all $x > 0$ and, second, that smaller λ and greater $Q(x)$ lead to a greater $R(-x)$.

First, decreasing λ has the same effect as increasing $Q(x)$ in Lemma 9 because both enter the density in (17) only through $\frac{Q(x)}{\lambda}$. Thus, decreasing λ leads to a likelihood-ratio shift downwards and by (17) to a greater $Q(x)$ for any $x > 0$.

Second, given $Q(x)$, $R(-x)$ is determined by the implementation-cost minimization problem. Inserting the constraint $x + R(x) + R(-x) = Q(x)$, we can rewrite the implementation cost Ψ with slight abuse of notation as a function of $R(-x), Q(x)$, and λ:

$$\Psi(R(-x), Q(x), \lambda) := \left(\int \frac{e^{\frac{\beta}{\lambda}Q(x)}}{1 + e^{\frac{\beta}{\lambda}Q(x)}} dF(\beta) \right) \psi(Q(x) - x + R(-x)) +$$

$$\left(1 - \int \frac{e^{\frac{\beta}{\lambda}Q(x)}}{1 + e^{\frac{\beta}{\lambda}Q(x)}} dF(\beta) \right) \psi(R(-x))$$

We show that the Ψ satisfies strictly increasing differences in $(R(-x), (Q(x), -\lambda))$, which implies that the solution $R(-x)$ to the minimization problem is strictly smaller, when $Q(x)$ is greater and λ is smaller by Topkis Theorem. By Topkis Characterization Theorem, it is sufficient to show that $\frac{d^2\Psi}{dQ(x)dR(-x)} > 0$ and $-\frac{d^2C}{d\lambda dR(-x)} = \frac{d^2C}{d(-\lambda)dR(-x)} > 0$. We can restrict attention to $Q(x) > 0$, $R(-x) < 0$ and $R(x) = Q(x) - x + R(-x) > 0$, which are satisfied at the optimum by Theorem 1 anyways.

$$\frac{d\Psi}{dR(-x)} = \left(\int \frac{e^{\frac{\beta}{\lambda}Q(x)}}{1 + e^{\frac{\beta}{\lambda}Q(x)}} dF(\beta) \right) \psi'(Q(x) - x + R(-x))$$

$$+ \left(1 - \int \frac{e^{\frac{\beta}{\lambda}Q(x)}}{1 + e^{\frac{\beta}{\lambda}Q(x)}} dF(\beta) \right) \psi'(R(-x))$$

$$\Rightarrow \frac{d^2\Psi}{dQ(x)dR(-x)} = \left(\int \frac{d}{dQ(x)} \frac{e^{\frac{\beta}{\lambda}Q(x)}}{1 + e^{\frac{\beta}{\lambda}Q(x)}} dF(\beta) \right) \psi'(Q(x) - x + R(-x))$$

$$+ \left(\int \frac{e^{\frac{\beta}{\lambda}Q(x)}}{1 + e^{\frac{\beta}{\lambda}Q(x)}} dF(\beta) \right) \psi''(Q(x) - x + R(-x))$$

$$- \left(\int \frac{d}{dQ(x)} \frac{e^{\frac{\beta}{\lambda}Q(x)}}{1 + e^{\frac{\beta}{\lambda}Q(x)}} dF(\beta) \right) \psi'(R(-x))$$

By differentiability of the integrand in Q and boundedness of this derivative, we exchanged integration over β and differentiation by $Q(x)$ according to Leibniz rule. Since ψ is strictly convex and minimal at zero, we have $\psi'(x) \gtrless 0 \Leftrightarrow x \gtrless 0$ and $\psi''(x) > 0$. By the constraints on $Q(x)$ and $R(-x)$, we have $R(-x) < 0$ and $Q(x) - x + R(-x) > 0$, so $\psi'(R(-x)) < 0$ and $\psi'(Q(x) - x + R(-x)) > 0$. Finally, by

$$\frac{e^{\frac{\beta}{\lambda}Q(x)}}{1 + e^{\frac{\beta}{\lambda}Q(x)}} > 0 \quad \text{and}$$

$$\frac{d}{dQ(x)} \frac{e^{\frac{\beta}{\lambda}Q(x)}}{1 + e^{\frac{\beta}{\lambda}Q(x)}} > 0$$

for $Q(x) > 0$, we obtain that $\frac{d^2\Psi}{dQ(x)dR(-x)}$ is strictly positive.

Regarding λ,

$$\frac{d^2\Psi}{d\lambda dR(-x)} = \left(\int \frac{d}{d\lambda} \frac{e^{\frac{\beta}{\lambda}Q(x)}}{1 + e^{\frac{\beta}{\lambda}Q(x)}} dF(\beta) \right) \psi'(Q(x) - x + R(-x))$$

$$- \left(\int \frac{d}{d\lambda} \frac{e^{\frac{\beta}{\lambda}Q(x)}}{1 + e^{\frac{\beta}{\lambda}Q(x)}} dF(\beta) \right) \psi'(R(-x))$$

By

$$\frac{d}{d\lambda} \frac{e^{\frac{\beta}{\lambda}Q(x)}}{1 + e^{\frac{\beta}{\lambda}Q(x)}} < 0$$

and we obtain that $-\frac{d^2\Psi}{d\lambda dR(-x)} < 0$ and hence $\frac{d^2\Psi}{d(-\lambda)dR(-x)} > 0$. $\qquad \square$

I   Proof of Theorem 3

Because the feedback structure reveals the material payoff by Assumption 4, the principal can choose the utility $u(z)$ for each message $z = (v, a, P) \in Z$. For $z = (v, a, P)$, we write $v(z)$ and $a(z)$ for the first and third component, respectively. The following Lemma uses of the fact that the optimal attention strategy under posterior-separable attention costs depends on the state-wise utility differences as shown in Whitmeyer and Zhang (2022).

Lemma 10. *A first-best attention strategy is induced only if for all elements z_1, z_2 of Z that are directly connected by an edge and have $a(z_1), a(z_2) \in S$,*

$$u(z_1) - u(z_2) = \frac{1}{\beta}(v(z_1) - v(z_2)).$$

Proof. By Theorem 4.3 in Whitmeyer and Zhang (2022), to induce an attention strategy that is optimal for the principal, the relative incentives, such as $u(a, \omega) - u(a_1, \omega)$ for $a \in S, \omega \in \Omega$, are uniquely pinned down. Since the agent discounts the relative incentives by β, for the relative incentives to equal those of the principal, they need to be scaled up by $\frac{1}{\beta}$. This is achieved if and only if for all $\omega \in \Omega$, $u(a, \omega) - u(a_1, \omega) = \frac{1}{\beta}(v(a, \omega) - v(a_1, \omega))$. This in turn is equivalent to all elements m_1, m_2 of Z that are directly connected by an edge (that is, are achieved at the same state) to satisfy $u(m_1) - u(m_2) = \frac{1}{\beta}(v(m_1) - v(m_2))$. $\square$

The above is an only-if statement. It ensures only that if we restricted the agent's choice to the consideration set S, a first-best attention strategy would be incentive-compatible (Whitmeyer and Zhang, 2022). To implement the first-best attention strategy when the whole decision problem A is feasible, it is sufficient to punish the actions in $A \setminus S$ with sufficiently negative emotions. As long as the principal deters the agent from choosing any action in $A \setminus S$, she is indifferent with regards to how strongly she punishes these actions as it does not affect her utility nor the implementation cost of realized emotions.

The utility difference being scaled up by $\frac{1}{\beta}$ is a transitive relation that extends to all nodes on the largest connected subgraph. If z and $\tilde{z}$ are on the same connected subgraph of Z, then there is a finite chain of elements $(z_1, ..., z_k)$ of Z such that $z = z_1$, $\tilde{z} = z_k$, and each adjacent pair is directly connected. If each of those directly connected pairs has a utility difference scaled up by $\frac{1}{\beta}$, then

$$u(\tilde{z}) - u(z) = (u(z_k) - u(z_{k-1})) + ... + (u(z_2) - u(z_1))$$
$$= \frac{1}{\beta}(v(z_k) - v(z_{k-1}) + ... + \frac{1}{\beta}(v(z_2) - v(z_1)) = \frac{1}{\beta}(v(\tilde{z}) - v(z)).$$

Conversely, if the utility difference of all pairs of nodes on the same connected subgraph is scaled up by $1/\beta$, then trivially the utility difference of directly connected nodes is scaled up by $1/\beta$.

Thus, an optimal attention strategy of the principal is implemented only if the utility difference of all elements on the same connected subgraph are scaled up by $1/\beta$. By Assumption 4, the sets of all messages that are consistent with some partition element $P \in \mathcal{P}$ are connected. By the definition of edges, such sets for different partition elements are not directly connected. In other words, Z has one connected subgraph for each partition element $P \in \mathcal{P}$. This pins down the utility function u except for one constant per subgraph, which is determined by the implementation-cost

minimization problem. More specifically, by the above, we can write

$$m(a, \omega) = \left(\frac{1}{\beta} - 1 \right) \left(v(a, \omega) - v(\mathcal{P}(\omega)) \right). \tag{22}$$

The remaining problem is to determine the "reference point" $v(\mathcal{P}(\omega))$ for each partition cell.

The optimal reference point is connected to the concept of central tendency from statistics. The central tendency $\mathcal{C}_f(X)$ of random variable X with respect to a strictly convex function f is defined as

$$\mathcal{C}_f(X) := \arg\min_c \mathbb{E}[f(X - c)],$$

which is known to be unique under strict convexity of f. Here, we consider the central tendency $\mathcal{C}_\psi$ with respect to the strictly convex implementation cost ψ.

Lemma 11. *Under implementation cost $\psi(x) = |x|^k$ with $k > 1$, the implementation cost minimizing solution has reference point $v(P)$ equal to the central tendency of payoff v on the partition element P with respect to the optimal attention strategy.*

Proof. First, the centrality measure of emotions must be zero, otherwise we could decrease the implementation cost. Let $P \in \mathcal{P}$. The attention strategy that is implemented defines a distribution over emotions. The implementation cost on this partition is the expectation of $\psi(u(a, \omega) - v(a, \omega))$. Under the optimal emotions, $\psi(u(a, \omega) - v(a, \omega) - c)$ is minimized at $c = 0$, otherwise one could further reduce the implementation cost. This implies by definition that the central tendency of emotions is zero.

Next, we show two properties of the central tendency. First, $\mathcal{C}_\psi(X + b) = \mathcal{C}_\psi(X) + b$ due to

$$\arg\min_c \mathbb{E}[\psi(X + b - c)] = b + \arg\min_c \mathbb{E}[\psi(X - c)].$$

Second, under $\psi(x) = |x|^k$ with $k > 1$, $\mathcal{C}_\psi(aX) = a\mathcal{C}_\psi(X)$ for all $a \in \mathbb{R}$. First of all, note that $\psi(ax) = |ax|^k = (|a||x|)^k = |a|^k |x|^k = |a|^k \psi(x)$. Then,

$$\arg\min_c \mathbb{E}[\psi(aX - c)] = a \arg\min_c \mathbb{E}[\psi(a(X - c))] = a \arg\min_c \mathbb{E}[|a|^k \psi(X - c)] =$$
$$a \arg\min_c \left(|a|^k \mathbb{E}[\psi(X - c)] \right) = a \arg\min_c \mathbb{E}[\psi(X - c)].$$

Finally, using the two previous facts we show that the reference point of P is the central tendency of the payoff v, with respect to the distribution induced by the implemented attention strategy on P. The reference point $v(\mathcal{P}(\omega))$ is the (hypothetical) payoff $\tilde{v}$ that obtains emotions zero according to equation (22). By equation (22), emotions are a linear function of the payoff $v(a, \omega)$. Then, the realized payoff is also a linear function f of emotions. The payoff $\tilde{v}$ is thus the image of 0 under f, $\tilde{v} = f(0)$. By the above, $\mathcal{C}_\psi(aX + b) = a\mathcal{C}_\psi(X) + b$, so $\mathcal{C}_\psi$ commutes with linear functions. By 0 being the central tendency of emotions, $\tilde{v}$ is the central tendency of the payoff. $\qquad\square$

As is known, the central tendency $C_\psi(X)$ of random variable X with respect to $\psi(x) = x^2$ equals the expectation of X. This concludes the proof.