

A Criterion of Model Decisiveness*

Jeffrey Yang[†]

First posted: April 8, 2023

Updated: June 17, 2025

Abstract

When faced with decision-relevant information, decision-makers are often exposed to a multiplicity of different models, or accounts of how information should be interpreted. This paper proposes a theory of model selection — an account of what models decision-makers find compelling, and ultimately adopt — based on the insight that individuals seek *decisive* models that provide clear guidance regarding the best course of action. The decisiveness criterion is characterized by a demand for extreme models, which generates inferential biases such as overprecision and confirmation bias, but predicts meaningful bounds on the extent of these biases. The dependence of the decisiveness criterion on the decision-maker’s objectives can produce documented patterns of preference reversals, rationalize seemingly contradictory patterns of inferential attribution errors, and generate novel predictions as to how belief polarization can arise along heterogeneity in decision-makers’ objectives. I discuss applications of the theory to financial decision-making, the provision of expert advice, and social learning through the exchange of models.

*I am indebted to Benjamin Enke, Matthew Rabin, Joshua Schwartzstein, and Tomasz Strzalecki for their excellent supervision and guidance. I also thank Peter Andre, Chiara Aina, Thomas Graeber, David Laibson, Shengwu Li, Gautam Rao, Andrei Shleifer, Andi Wang and Harvard PhD workshop participants for helpful comments and suggestions.

[†]Department of Economics, Harvard University. jeffrey_yang@g.harvard.edu.

1 Introduction

When faced with uncertainty over the correct course of action, decision-makers are often confronted with an abundance of potentially decision-relevant information. It is increasingly understood that in order to make sense of this information, decision-makers often reason through the lens of a model – an account of how information should be interpreted. When an investor decides to sell a security because of a recent price drop, she acts based on a model that relates price patterns to future performance. When a manager rejects a candidate on the basis of a botched interview in spite of an impressive track record, she appeals to a model specifying what information is most diagnostic of candidate productivity. A voter who bases her political views on the reporting of certain news outlets while remaining skeptical toward others operates according to a model that specifies the credibility of different sources. In each of these cases, the decision-maker is guided by a framework, whether implicit or explicit, that specifies what information is relevant to the decision at hand, and what conclusions should be drawn from that information.

Decision-makers are often exposed to multiple candidate models: there are many plausible ways to value a firm, evaluate a job candidate, or interpret the news, and these models may be supplied by experts, encountered through social exchange, or developed through introspection. In contrast to the rational expectations view, under which decision-makers interpret data according to a well-calibrated model, decision-makers may not know the true model that should guide their inference, and so must decide which model to adopt for a given situation. In many cases, decision-makers are drawn to models that lead them astray: professional forecasters overreact to market signals, hiring managers put too much stock in noisy interview measures, and voters are misled by baseless claims and misinformation. When individuals seek to interpret information, what kinds of models do they find compelling, and ultimately use to guide their decision-making?

One potential answer stems from the basic insight that people dislike indecision, an observation which draws support from research in psychology. First, research on cognitive dissonance has demonstrated that individuals seek to avoid the psychological discomfort that arises from holding cognitions that are inconsistent with each other (Festinger, 1957), and modern accounts of dissonance theory emphasize the role between dissonance and indecision, wherein dissonance is aroused when “cognitions with action implications are in conflict with each other, making it difficult to act” (Harmon-Jones et al., 2015). A second line of research studies individuals’ need for cognitive closure, conceptualized as a “desire to have a definite answer to a question, as opposed to uncertainty, confusion and ambiguity” (Kruglanski and Fishman, 2009). Under this account, individuals desire closure because it affords “a base for action”, and research has studied how the need for closure affects how individuals process information and revise their beliefs.¹

If individuals are indeed averse to indecision, they may seek models that are *decisive*: that provide clear guidance regarding the best course of action. In this paper, I propose

¹More recently, Proulx and Inzlicht (2012) proposes a psychological framework of sense-making to unify these two accounts, in addition to a broad range of related accounts and evidence in psychology. Under this account, events that violate *meaning* — defined as a set of expected relationships that serve as a “guide for action” — leads to a physiological state of aversive arousal, which individuals seek to reduce.

a formal notion of decisiveness, under which a model is decisive to the extent it reduces the decision-maker’s residual uncertainty over the optimal course of action; I then analyze the inferences and choices of a decision-maker who is drawn to such models in interpreting information. In my formal framework, a decision-maker chooses from a menu of actions A , and observes the realization of data s that may be informative about the payoff-relevant state θ . To take an example in which the decision-maker is evaluating a job candidate, A may be the choice of whether to hire or reject the candidate, s is the information about the candidate observed by the decision-maker, such as their qualifications and interview performance, and θ is the underlying productivity of the candidate. In interpreting this data, the decision-maker does not have access to the true data-generating process but instead entertains a set of models: likelihood functions that map data to posterior beliefs over states. For example, the decision-maker may entertain models under which interview performance is highly diagnostic of productivity, in addition to models that instead emphasize the diagnostic value of the candidate’s track record.

A model is decisive to the extent it recommends an action that, *under the beliefs induced by the model*, is close to the ex-post optimal course of action. That is, given $p^m(\theta|s)$, the posterior over states θ induced by model m for the observed data s , the decisiveness of m is

$$I(m|s) = - \left[\sum_{\theta} \max_{a' \in A} u(a', \theta) p^m(\theta|s) - \max_{a \in A} \sum_{\theta} u(a, \theta) p^m(\theta|s) \right]$$

where $u(a, \theta)$ gives the payoff of action a if the state is θ . Note that decisiveness corresponds to a payoff-metric measure of residual uncertainty: $-I(m|s)$ is equal to the decision-maker’s willingness to pay to resolve remaining uncertainty about the state, under the beliefs induced by m . In other words, a model is decisive if a decision-maker operating under that model ascribes little value — if such an alternative was available to her — to resolving uncertainty prior to making her decision, instead of choosing the course of action recommended by the model.

Decisive models make strong recommendations, but they may also lead the decision-maker astray: a model suggesting that stock returns are highly predictable on the basis of past returns will be decisive, but is also far from the truth. Decisive models need not coincide with the true model governing the data-generating process, and to the extent the decision-maker entertains such models, the decisiveness criterion will distort their beliefs and decision-making. Importantly, in my framework the decision-maker does not know the true model — she does not willingly commit to a model she knows is wrong. Instead, my framework corresponds to a cognitive process in which the DM entertains a set of candidate models and tries each on “for size”, assessing each model according to its implications for her decision and ultimately adopting the model that feels the most compelling. As such, this framework departs from certain theories of motivated beliefs in which the decision-maker seeks to maintain a set of beliefs in spite of the errors that may arise when she acts on them; rather, the decision-maker in my framework seeks decisive models precisely because she is motivated to identify the correct course of action.

This paper analyzes the systematic distortions in inference and choice produced by model selection under the decisiveness criterion. In particular, I show that the decisive-

ness criterion is characterized by a demand for models that render the decision-maker's environment more predictable. This causes the decision-maker to exhibit forms of both overprecision and confirmation bias: when initially uncertain over the optimal course of action, the decision-maker favors models that overstate the diagnostic value of the data, and when sufficiently predisposed towards a course of action, the decision-maker favors models that minimize the diagnostic value of disconfirmatory data. I also show how the dependence of the decisiveness criterion on the decision-maker's objectives can produce documented patterns of preference reversals, rationalize seemingly contradictory patterns in social attribution, and generate novel predictions as to how belief polarization can arise along heterogeneity in decision-makers' objectives. Finally, I demonstrate how choice under the decisiveness criterion exhibits an aversion to various forms of hedging against uncertainty.

To illustrate the basic implications of the decisiveness criterion, consider a simple example. Suppose that a manager is deciding whether or not to make an offer to a potential internal hire to join her team. The candidate is either high-productivity (θ_h) or low-productivity (θ_l), and the manager believes either possibility is equally likely ex-ante. The manager's payoffs are as follows:

	θ_h	θ_l
<i>hire</i>	v	$-k$
<i>reject</i>	0	0

That is, the manager wishes to hire the candidate only if they are high-productivity.

The manager has access to two components of information on the candidate: the candidate's track record in their previous role, s^R , and the candidate's performance in a newly developed interview assessment tailored to the current role s^I . Suppose that performance in either component $c \in \{R, I\}$ can either be high ($s^c = 1$) or low ($s^c = 0$), and that the manager is uncertain over the diagnostic value of each piece of information. In particular, suppose the manager entertains three models: m_R , under which only the candidate's track record has diagnostic value, m_I , under which only the interview assessment has diagnostic value, and m_{RI} , under which both components hold equal diagnostic value. In particular, denoting $L_m(s^R, s^I) = \frac{m(s^R, s^I | \theta_h)}{m(s^R, s^I | \theta_l)}$ as the likelihood ratio associated with model m , we have

$$\begin{aligned}
L_{m_R}(s^R, s^I) &= \begin{cases} 4 & s^R = 1 \\ 1/4 & s^R = 0 \end{cases} \\
L_{m_I}(s^R, s^I) &= \begin{cases} 4 & s^I = 1 \\ 1/4 & s^I = 0 \end{cases} \\
L_{m_{RI}}(s^R, s^I) &= \begin{cases} 4 & s^I = s^R = 1 \\ 1 & s^I \neq s^R \\ 1/4 & s^I = s^R = 0 \end{cases}
\end{aligned}$$

Suppose that the candidate's track record demonstrates high performance in their previous role, but that they perform poorly in the interview assessment: $(s^R, s^I) = (1, 0)$.

Which model does the manager adopt under the decisiveness criterion? Below, we consider how the manager acts under two different hiring regimes, each of which dictate the relative costs and benefits of hiring.

Growth regime: $v = 4, k = 1$. In this regime, the benefits of hiring a high-productivity candidate are greater than the costs of hiring a low-productivity candidate; the manager is therefore predisposed to hiring the candidate. Here, m_R is the most decisive model. The intuition is as follows: under m_R the data suggests that the candidate is likely high productivity, providing the manager with yet greater justification for hiring the candidate. On the other hand, under m_{RI} the data are inconclusive about candidate productivity, and so this model provides weaker justification for hiring compared to m_R . Under m_I , the data are bad news about the productivity of the candidate, an interpretation which, given the manager’s decision problem, would result in her being maximally uncertain over whether to hire or reject the candidate. As such, m_R provides the most decisive recommendation of the three models.

Downsizing regime: $v = 1, k = 4$. In this regime, hiring costs are high, and so the manager is predisposed to rejecting the candidate. Here, m_I is the most decisive model. The intuition is analogous to the case above: while both m_I and m_{RI} recommend that the manager reject the candidate, m_I provides a stronger justification toward this course of action, whereas under m_R , the manager is maximally uncertain over whether to hire or reject.

This example illustrates two key properties of the decisiveness criterion. First, the criterion tends to privilege extreme models — formally, models that cannot be expressed as a mixture of other models the DM entertains. Note that in either case, the model m_{RI} , which ascribes some diagnostic value to both the candidate’s track record and interview performance, and can be expressed as the mixture of m_R and m_I , is never selected. The intuition is as follows: if the manager prefers *reject* under m_{RI} , m_I provides yet stronger justification towards *reject*, whereas if the manager prefers *hire* under m_{RI} , m_R provides yet stronger justification towards that course of action; the manager must find m_{RI} less decisive than one of the extreme models, which will provide greater certainty regarding her optimal hiring decision. Second, the decisiveness criterion depends crucially on the decision-maker’s objectives, as the difference between the two regimes demonstrate: in each regime, the model that supports the action that the manager is predisposed towards choosing is more decisive than the model that provides disconfirmatory evidence.

The remainder of the paper is organized as follows: Section 2 develops the general framework and introduces the decisiveness criterion. Section 3 characterizes the implications of the decisiveness criterion for model selection, holding fixed the action space and payoffs. I show how model selection under the decisiveness criterion is principally characterized by two key conditions, which formalize the notion in which the criterion privileges models that render the decision-maker’s environment predictable: 1) the decision-maker has a tendency to adopt extreme models — those that cannot be expressed as a mixture of other models she entertains, and 2) the decision-maker has a preference for models that induce high certainty in a single state. I apply these results to show how the decisiveness

criterion generates forms of two documented patterns of biased inference: overprecision and confirmation bias. In particular, the criterion predicts that if the decision-maker has greater initial uncertainty with respect to her available courses of action, she will tend to exhibit overprecision, seeking models that overstate the informational content of her signals. However, as the decision-maker’s prior sufficiently favors one course of action, she will no longer uniformly exhibit overprecision, but instead seek models that explain away disconfirmatory information, producing a form of confirmation bias. Finally, I illustrate the relationship between my model selection criterion and related criteria, such as the Blackwell ordering on experiments and notions of model parsimony.

Section 4 studies how model selection under the decisiveness criterion varies with the payoffs and objectives of the decision-maker. First, I demonstrate that the decisiveness criterion generates a “sour grapes” effect: the addition of an unchosen action a will lead the decision-maker towards models that ascribe a low value to choosing a , as well as similar actions. I show how this sour grapes effect predicts context effects documented in Tversky and Shafir (1992), who show how the addition of an action to the choice set can reduce subjects’ propensity to choose similar, competing actions, and instead induce them to choose dissimilar actions. Next, I show how under the decisiveness criterion, the attractiveness of a model is increasing in the relative attractiveness of the actions that the model recommends. I demonstrate how this property can account for seemingly contradictory patterns in social attribution in which individuals neglect the confounding influence of situational factors when inferring dispositional traits of others from their behavior, yet make the opposite error when interpreting poor behavior from ingroup members (Vonk and Konst, 1998). In another application, I show how the same force can generate belief polarization resulting from heterogeneity in decision-makers’ objectives and preferences.

In Section 5, I discuss implications for choice. I show that choice under the decisiveness criterion is characterized by an aversion to hedging or diversification. In addition, I study comparative statics of choice under the decisiveness criterion, focusing on a notion of relative aversion to *C-diversified* actions — that is, actions whose payoffs fall short of the maximal payoff that can be achieved in each state by a constant. I show that the decision-maker will be more diversification-averse if she entertains a larger set of models, and that conversely, a greater level diversification-aversion reveals that the decision-maker entertains a larger set of models. In Appendix A.5, I provide an axiomatic characterization of the model and discuss its identification properties. These results parallel those in Stoye (2011), which studies a closely related model of min-max regret.

Section 6 discusses additional applications of the decisiveness criterion. Whereas the basic framework takes the set of models the decision-maker entertains as primitive, each of these applications considers a process that may give rise to this set. In the first application, I analyze a setting in which expert advisors supply models to the decision-maker, and apply the decisiveness criterion to shed light on why individuals are drawn to certain advice. I demonstrate that this force, in conjunction with competition between advisors, induces advisors to inflate the certainty of their recommendations. In the second application, I analyze a setting in which social learning occurs through the exchange of models, following the framework introduced in Schwartzstein and Sunderam (2022), and demonstrate that the decisiveness criterion predicts group polarization arising from such a process.

Section 7 extends the theory by specifying how the set of models the decision-maker entertains may be constrained by their plausibility. I discuss two plausibility constraints, one which assumes a cost of adopting far-fetched beliefs, the other which stipulates an entry condition on the models the DM can entertain based on the notion of model fit developed in Schwartzstein and Sunderam (2021).

Related Literature

This paper directly contributes to an active literature studying the implications of model selection, which has focused on several distinct model selection criteria. One criterion is model fit, corresponding to the notion that decision-makers are drawn to models that explain the data well. In particular, the formal framework employed in this paper draws directly from Schwartzstein and Sunderam (2021), who analyze “model persuasion” — in which a persuader influences the model the receiver uses to interpret the realized data — and assume that receivers employ a model selection criteria based on fit. Aina (2022) builds on this framework, analyzing a setting in which the persuader commits to a set of models prior to the realization of the data, whereas Schwartzstein and Sunderam (2022) adapt the framework to study social learning based on the exchange of models. In contrast, this paper focuses primarily on non-strategic contexts, in which the set of models the decision-maker entertains is primitive; however, Section 6 discusses applications of the decisiveness criterion to strategic and social learning contexts. Operating under different formal frameworks, Izzo et al. (2021) and Hong et al. (2007) study the implications of fit-based model selection, and focus on political persuasion and a non-strategic financial decision-making context, respectively. Another criterion that has received attention in the literature is optimism, corresponding to the notion that decision-makers are drawn to “hopeful narratives”. Eliaz and Spiegler (2020) formalize narratives as causal models (directed acyclic graphs) to study political persuasion, and posits a model selection criteria based on optimism, whereas Caplin and Leahy (2019) studies the implications of model selection based on optimism for belief polarization, trading decisions, and the formation of asset bubbles.

The decisiveness criterion provides an account of what makes a model compelling that is complementary to these existing approaches, and in particular delivers predictions distinct from those of selection criteria based on fit or optimism. The fit criterion, for one, only concerns how well models can explain past data, and is silent on how individuals may be drawn to models due to their implications for future action. For example, one might imagine that employers continue to place high stock in unstructured interviews (Dana et al., 2013) not necessarily because such interviews have demonstrated explanatory power over the productivity of past hires, but because taken at face value, they provide strong recommendations for hiring decisions. Similarly, investors may find technical analysis appealing not just because it provides an explanation for realized price trends, but also because it often provides a clear recommendation on whether to buy or sell. Importantly, because a fit-based model selection criterion is backward-looking, without further structure and assumptions, such a criterion cannot account for the possibility that the model a decision-maker adopts can be shaped by their objectives. On the other hand, while an optimism-based criterion

takes seriously the idea that a model’s forward-looking implications may matter for model selection, it is clear that decision-makers also regularly adopt models that do not induce optimistic beliefs. For example, employers may overstate the diagnostic value of an interview even if it produces bad news about the productivity of the candidate, and voters are swayed by political spin often solely focused on casting the opposition in a negative light. Because a decisive model need not lead to optimistic beliefs, as illustrated in the motivating example, the decisiveness criterion can rationalize the adoption of such models.

Though the decisiveness criterion aims to capture a psychological notion distinct from regret — namely, that individuals are drawn to clear recommendations that resolve indecision — the criterion is formally equivalent to a formulation of expected regret, and therefore relates to a large literature studying the implications of regret aversion in decision-making (e.g. Bell, 1982; Loomes and Sugden, 1982; Hayashi, 2008; Sarver, 2008; Stoye, 2011). This literature studies how regret aversion may affect choices over actions or menus of actions, corresponding to the psychology that the desire to avoid regret may affect choices prior to the realization of uncertainty. This paper, on the other hand, focuses on the implications of a regret-based criterion for model selection and its subsequent implications for decision-making, motivated by the notion that individuals find models compelling to the extent that they guide future action. Formally, this theory has a tight connection to a model of multiple-priors min-max regret studied in Hayashi (2008) and subsequently Stoye (2011), in which the decision-maker chooses the action to maximize expected utility (minimize expected regret) under the regret-maximizing prior belief. My theory corresponds to a model in which the decision-maker instead chooses to maximize expected utility under the regret-minimizing prior belief. While the two models produce contrasting predictions for both model selection and choice — for instance, min-max regret would predict that the DM always selects m_{RI} in the hiring example analyzed above, whereas the decisiveness criterion predicts that the DM never selects m_{RI} — the axiomatic characterization of my model parallels that of Stoye (2011).

One paper that takes a related conceptual approach is Eyster et al. (2021), which develops a model of ex-post rationalization in which a decision-maker adopts preferences to minimize an *ex-post* notion of regret with respect to past choices, motivated by evidence for the sunk-cost fallacy. This approach is both formally distinct from and complementary to the approach taken in this paper, which develops a theory in which the decision-maker adopts models that minimize *ex-ante* regret. In the hiring example analyzed above, for instance, model selection under ex-post rationalization depends crucially on each model’s implications for the optimality of actions the decision-maker has already taken, such as previous hiring decisions. On the other hand, model selection under the decisiveness criterion depends only on each model’s implications for the decision problem the decision-maker currently faces.

2 Framework and Decisiveness Criterion

2.1 The Decisiveness Criterion

Consider a decision-maker (the DM) with full-support priors q over a finite set of payoff relevant states Θ , and updates her beliefs about Θ given a signal $s \in S$, where $|S| \geq 2$.² To relate observed signals to states, the DM relies on a model $m : \Theta \rightarrow \Delta(S)$, a likelihood function mapping states to probability distributions over signals. The DM entertains a closed set of models $M \subseteq \mathcal{M}$, where \mathcal{M} denotes the set of all models. Let m_T denote the true model that governs the data-generating process.

The DM faces a *decision problem* $\mathcal{D} = (A, u)$, where A is a set of actions and $u : A \times \Theta \rightarrow \mathbb{R}$ is the payoff function. Assume that \mathcal{D} is well defined, in that for any beliefs there will be a set of expected-utility maximizing actions, i.e. $\arg \max_{a \in A} \sum_{\theta} u(a, \theta) p(\theta)$ is nonempty for all $p \in \Delta(\Theta)$.

The timing is as follows: nature first draws the state θ and the signal s is generated according to m_T and is observed by the DM. The DM then adopts a model $m \in M$, and forms posteriors $p^m(\cdot|s)$ according to Bayes rule, where $p^m(\theta|s) = m(s|\theta)q(\theta)/(\sum_{\theta'} m(s|\theta')q(\theta'))$. The DM then takes the action that maximizes the model-induced posterior expected utility; let

$$A_{\mathcal{D}}^m(s) \equiv \arg \max_{a \in A} \sum_{\theta} u(a, \theta) p^m(\theta|s)$$

denote the set of actions recommended by m , and say that m recommends a from \mathcal{D} if $a \in A_{\mathcal{D}}^m(s)$. After the action is taken, the DM's payoffs are then realized according to u and the realized state θ .

We now specify the criterion that governs the DM's model selection. For a posterior belief p and decision problem \mathcal{D} , let

$$R_{\mathcal{D}}(p) = \sum_{\theta} \max_{a'} u(a', \theta) p(\theta) - \max_a \sum_{\theta} u(a, \theta) p(\theta) \quad (1)$$

denote the *residual uncertainty* associated with p : it is the average difference in utility between the ex-post optimal action in each state and the ex-ante optimal action under p . $R_{\mathcal{D}}(p)$ can be interpreted as the DM's willingness-to-pay to resolve uncertainty prior to choosing an action given beliefs p . Define the *decisiveness* of a model m given realized signal s as

$$I_{\mathcal{D}}(m|s) = -R_{\mathcal{D}}(p^m(\cdot|s)). \quad (2)$$

The DM adopts a model that maximizes decisiveness, satisfying $m \in \arg \max_{m \in M} I_{\mathcal{D}}(m|s)$. Decisive models make strong prescriptions about the optimal course of action, where here a prescription is strong in the sense that it recommends an action that is likely close to the ex-post optimal action. That is, a model is decisive to the extent the DM ascribes little

²The assumption that q is full-support is without loss, as we can restrict attention to states on which the DM's prior places positive probability.

value — if such a course of action was available — to resolving uncertainty prior to making her decision, instead of choosing the course of action recommended by the model.

The decisiveness criterion captures a notion of what it means for a model to guide action. Two related notions of guidance that are not captured by this criterion are model consistency and parsimony. Individuals may seek models that provide *consistent* guidance — that is, models that recommend the same action under many realizations of the data. For example, consider a model of stock returns under which the price path follows a random walk with positive drift; for any history of returns, the model’s recommendation is the same: the DM should buy and hold. Such a model produces a consistent recommendation, in contrast to models that recommend various strategies to time the market depending on past returns. The decisiveness criterion will tend to privilege models of the latter kind over the former: under this criterion, it is not the consistency of guidance that matters, but rather its strength as measured by residual decision uncertainty. Individuals may also be drawn to *parsimony* — that is, they may see a model as providing better guidance to the extent it provides a simplified way of translating data into action, for instance, by directing focus to certain features of the data while ignoring others. In Section 3.3, I show that the decisiveness criterion does not universally privilege such simplified models, but rather generates predictions as to when such models will be adopted.

Note that in this framework, model selection occurs *ex-post*: after the signal is realized, the DM evaluates each model based on its decisiveness, given the signal realization, and adopts the most decisive model. One might instead imagine an account in which the DM evaluates each model according to an *ex-ante* notion of decisiveness, and adopts a model prior to the signal realization. The *ex-post* notion may be more appropriate for analyzing situations in which the DM faces a decision problem with an idiosyncratic information structure, as in the motivating example, in which the manager seeks to make a one-off hire on the basis of information specific to that particular candidate. Furthermore, because the *ex-post* decisiveness of a model can be evaluated solely on the basis of its recommendation given the observed signal realization, whereas *ex-ante* decisiveness requires the DM to consider what the model would recommend for all possible signal realizations, *ex-post* decisiveness may be more reasonable criterion for studying settings in which the set of possible signal realizations may be large, unknown, or otherwise hard to imagine. The *ex-ante* perspective, on the other hand, may be more appropriate for situations in which the DM faces a sequence of decision problems that share a common information structure — for example, if the hiring manager instead sought to evaluate a population of candidates on the basis of their performance in a standardized test. While this paper focuses on the *ex-post* notion of decisiveness, Appendix A.4 develops a formulation of *ex-ante* decisiveness and compares these two formulations.

Finally, the fact that the decisiveness criterion is defined with respect to a decision problem raises the question of how the theory can be applied to situations where the DM does not face a particular decision problem but nevertheless has the opportunity to learn from data. In such situations, one natural approach is to assume that the DM learns from data as if they face a prediction problem — that is, as if they are tasked with reporting a set of beliefs in $A = \Delta(\Theta)$ and paid according to a scoring function $u : \Delta(\Theta) \times \Theta \rightarrow \mathbb{R}$. In Appendix A.1, I characterize model selection under the class of prediction problems with

incentive-compatible scoring rules.

2.2 Examples

The following set of examples illustrate the basic mechanics of the decisiveness criterion. In each of these examples, we take the realized signal as given and work directly with model-implied posteriors $p^m(\cdot|s)$.

Example (Hiring Decision). Consider the hiring decision from the introduction. Here, the payoff relevant states are the productivity of the candidate, $\Theta = \{\theta_l, \theta_h\}$, and the DM's decision problem is given by the actions $A = \{hire, reject\}$ and the payoff function

$$u(a, \theta) = \begin{cases} v & a = hire, \theta = \theta_h \\ -k & a = hire, \theta = \theta_l \\ 0 & a = reject. \end{cases}$$

Fix a model m , and let $p_h^m = p^m(\theta_h|s)$ denote the model-implied posterior belief in θ_h . For $p_h^m \geq \frac{k}{v+k}$ the DM chooses *hire*, and chooses *reject* otherwise. The residual uncertainty associated with m is then

$$R_{\mathcal{D}}(p^m(\cdot|s)) = \begin{cases} k(1 - p_h^m) & p_h^m \geq \frac{k}{v+k} \\ vp_h^m & \text{otherwise.} \end{cases}$$

That is, models that induce beliefs toward either extreme tend to be more decisive, and the range for which decisiveness is increasing in p_h^m is increasing in the value of hiring a high-productivity worker v , and decreasing in the cost of hiring a low-productivity worker k .

Example (Prediction Market). The DM participates in a prediction market based on a binary state $\Theta = \{\theta_A, \theta_B\}$. There are two assets, A and B , which pay 1 if the corresponding state is realized and 0 otherwise. The assets are priced at π_A and π_B , respectively, with $\pi_A + \pi_B = 1$.

The DM is endowed with wealth w and chooses how much of each asset to purchase, (x_A, x_B) , and is risk averse with log utility. It is never strictly optimal for the DM to hold positive amounts of each asset, and so letting $a = x_A - x_B$, it is without loss to describe the DM's decision problem using the action space $A = [-w/\pi_B, w/\pi_A]$ and payoff function

$$u(a, \theta) = \begin{cases} \ln(w + a(1 - \pi_A)) & \theta = \theta_A \\ \ln(w - a\pi_A) & \theta = \theta_B. \end{cases}$$

Fix a model m , and let $p_A^m = p^m(\theta_A|s)$ denote the model-induced posterior belief in θ_A . First-order conditions imply that under m , the optimal allocation is $a = \frac{(p_A^m - \pi_A)}{(1 - \pi_A)\pi_A} w$.

Note that under state θ_A , the DM's ex-post optimal choice is to invest all of her wealth into asset A , which yields a payoff of $\ln(w/\pi_A)$, and under state θ_B , the DM's ex-post

optimal choice is to invest all of her wealth into asset B , which yields a payoff of $\ln(w/(1 - \pi_A))$. The residual uncertainty associated with m is then

$$R_{\mathcal{D}}(p^m(\cdot|s)) = -\ln(1 - p_A^m)(1 - p_A^m) - \ln(p_A^m)p_A^m$$

That is, the decisiveness criterion predicts that for this decision problem, the agent will select the model that minimizes posterior entropy, given the data.

3 Properties of Selection Under Decisiveness

In this section, I show that model selection under the decisiveness criterion is principally characterized by a preference for extreme models — those that cannot be expressed as a mixture of other models she entertains. In a series of applications, I discuss how this extremeness property can predict both existing and novel forms of overprecision and confirmation bias. I then discuss the relationship between the decisiveness criterion and related model selection criteria.

3.1 Characterization of Model Selection Rule

For a decision problem \mathcal{D} and signal s , let $C_{\mathcal{D}}(M|s) \equiv \arg \max_{m \in M} I_{\mathcal{D}}(m|s)$ denote the model choice correspondence under the decisiveness criterion. Because the regret function is concave in beliefs for any decision problem, $C_{\mathcal{D}}$ satisfies the following extremeness property:

Proposition 1. *For $m, m' \in M$, if $m, m' \notin C_{\mathcal{D}}(M|s)$, then for any $\lambda \in (0, 1)$, $\lambda m + (1 - \lambda)m' \notin C_{\mathcal{D}}(M|s)$. Furthermore, if $A_{\mathcal{D}}^m(s)$ and $A_{\mathcal{D}}^{m'}(s)$ are disjoint, then $\lambda m + (1 - \lambda)m' \notin C_{\mathcal{D}}(M|s)$.*

In words, under the decisiveness criterion, the DM tends to adopt extreme models — models that cannot be expressed as a mixture of other models in M — because the decisiveness of a composite model obtained by averaging two models cannot exceed the decisiveness of both of those models. Furthermore, if those two models recommend different actions, then the composite model delivers strictly lower decisiveness.

An immediate consequence of Proposition 1 is that the decisiveness criterion “biases” the DM against model averaging: if the true model m_T is an average over models in M , Proposition 1 implies that m_T will never deliver the strictly highest decisiveness among the models $M \cup \{m_T\}$, and that so long as at least two models in M recommend different actions, the DM will never select the true model.

For intuition as to why the DM finds such composite models unattractive, consider the following example: the DM is trying to predict whether it will rain in an hour, and entertains two models. m_1 says that current humidity is the key predictor of rain, and m_2 says that cloud cover is the key predictor. Now consider a composite model $\bar{m} = \frac{1}{2}m_1 + \frac{1}{2}m_2$ that says that some mixture of humidity and cloud cover predict rain — the composite model can be interpreted as saying that on any given day, humidity is the key predictor of rain with 50% probability, and otherwise cloud cover is the key predictor of rain. Both m_1 and m_2 provide a specific lens through which to interpret the data, whereas \bar{m} leaves the DM

uncertain over which of these two interpretations to use. As such, the DM will find that one of m_1 or m_2 provides more compelling guidance than the composite model \bar{m} .

Other than the extremeness property, what other properties govern model adoption under the decisiveness criterion? The proposition below states that under regularity conditions, the extremeness property, in conjunction with another property stating that models that resolve all uncertainty must be adopted, fully characterize $C_{\mathcal{D}}$.

Proposition 2. *A model choice correspondence C satisfies*

1. *Sen's α, β : If $m \in M \subseteq M'$ and $m \in C(M')$, then $m \in C(M)$. Also, if $m, m' \in C(M)$, $M \subseteq M'$ and $m' \in C(M')$ then $m \in C(M')$.*
2. *Continuity: For all $m \in \mathcal{M}$, $\{m' \in \mathcal{M} : m' \in C(\{m, m'\})\}$ and $\{m' \in \mathcal{M} : m \in C(\{m, m'\})\}$ are closed.*
3. *Scale Invariance: For any $m \in \mathcal{M}$, if m' satisfies $m'(s|\theta) = \lambda m(s|\theta) \forall \theta, \lambda > 0$, then $C(\{m, m'\}) = \{m, m'\}$.*
4. *Extremeness: For $m, m' \in M$, if $m, m' \notin C(M)$, then for any $\lambda \in (0, 1)$, $\lambda m + (1 - \lambda)m' \notin C(M)$.*
5. *Certainty Preference: If $m(s|\theta) = 1$ for any $\theta \in \Theta$ and $m \in M$, then $m \in C(M)$.*

if and only if there exists a decision problem \mathcal{D} and signal s such that $C(M) = C_{\mathcal{D}}(M|s)$ for all $M \subseteq \mathcal{M}$.

Sen's α, β reflects the fact that the decisiveness criterion induces a total order over \mathcal{M} for a given decision problem. Scale Invariance ensures that models that induce the same posterior over states are equivalent under the model selection criteria.

Extremeness captures the property discussed in Proposition 1. Certainty Preference says that any model that induces certainty in an outcome must always be adopted, as such a model entirely eliminates residual uncertainty. Together, Extremeness and Certainty Preference formalize the notion in which the decisiveness criterion privileges models that render the decision-maker's environment predictable. Below, I show how these properties can account for documented patterns of belief updating, as well as produce new predictions.

3.2 Applications

3.2.1 Overprecision

A large experimental literature has documented that individuals exhibit overprecision — the tendency to be excessively confident in the accuracy of one's beliefs (Moore et al., 2015). In a common experimental paradigm, subjects are asked to predict a binary outcome, and state their belief that their prediction is correct. Across a variety of domains, subjects' estimates of their own accuracy exceed their actual hit rates. Similarly, studies eliciting subjects' subjective confidence intervals around their estimates of continuous outcomes find that these confidence intervals tend to be too narrow.

A natural consequence of Proposition 1 is that the DM will exhibit overprecision in the following sense — the DM will tend to overstate the informational content of her signal. Formally, fix any decision problem \mathcal{D} and the realization of the signal s . Let m_T denote the true model that describes the DM’s actual experiment, and let m_\emptyset denote an uninformative model for which $m(s|\theta)$ is constant in θ . Proposition 1 implies the following result:

Corollary 1. *Consider a model m satisfying $\lambda m + (1 - \lambda)m_\emptyset = m_T$ for $\lambda \in (0, 1)$. If $I_{\mathcal{D}}(m_T|s) > I_{\mathcal{D}}(m_\emptyset|s)$, then $I_{\mathcal{D}}(m|s) > I_{\mathcal{D}}(m_T|s)$.*

Here, m can be interpreted as a model that overstates the informational content of the true experiment m_T , which is a mixture of m and noise. Corollary 1 states that to the extent the DM entertains such models that overstate the informativeness of the their experiment, those models will be favored over the true model, so long as the DM finds the true model more decisive than her prior.

Example (Overprecision in Hiring). Suppose a hiring manager wishes to learn whether a candidate is high-productivity (θ_h) or low productivity (θ_l), and holds a prior belief q . The employer observes a signal s of productivity and chooses from a set of actions (e.g. hiring/rejecting the candidate). Suppose that $I_{\mathcal{D}}(m_T|s) > I_{\mathcal{D}}(m_\emptyset|s)$ — that is, the manager’s signal reduced her residual uncertainty regarding which action to take under the true model. Suppose the manager also entertains a model m that overstates the decisiveness of her signal in the sense described earlier: m satisfies $\lambda m + (1 - \lambda)m_\emptyset = m_T$ for $\lambda \in (0, 1)$; Corollary 1 implies that the hiring manager will adopt m over the true model m_T . Given the binary state setting, the implications for the manager’s posterior beliefs are straightforward: if under the true model, s is good news about productivity, with $q(\theta_h) < p^{m_T}(\theta_h|s)$, then m induces yet a higher belief in productivity, with $p^m(\theta_h|s) > p^{m_T}(\theta_h|s)$; similarly if under the true model s is bad news about productivity, then m induces yet a lower belief in productivity. In other words, the manager overstates her ability to discern high-productivity candidates from low-productivity candidates.

To relate this example to the two-alternative forced-choice paradigm used in experimental tests for overprecision, suppose that the manager’s decision problem consists of a binary prediction $g \in \{\theta_l, \theta_h\}$ regarding the candidate’s productivity, as well as an estimate ϕ of the probability that the prediction will be correct. Formally, we have $A = \{\theta_h, \theta_l\} \times [0, 1]$, with $u((g, \phi), \theta) = u_1(g, \theta) + u_2((g, \phi), \theta)$, where

$$u_1(g, \theta) = \begin{cases} 1 & g = \theta \\ 0 & \text{otherwise} \end{cases}$$

$$u_2((g, \phi), \theta) = \begin{cases} -(1 - \phi)^2 & g = \theta \\ -\phi^2 & \text{otherwise} \end{cases}$$

That is, the manager is incentivized to predict the candidate’s type, and is incentivized to provide an accurate assessment of how likely her prediction is to be correct.

Suppose that $q(\theta_h) = 0.5$, and that as before, the DM entertains the true model m_T , as well as the model m as defined above. Here, we have $I_{\mathcal{D}}(m_T|s) \geq I_{\mathcal{D}}(m_\emptyset|s)$ for all s , which implies that $I_{\mathcal{D}}(m|s) \geq I_{\mathcal{D}}(m_\emptyset|s)$: the manager overreacts to her signal about productivity.

Consider the case where s is good news about productivity — the manager predicts productivity to be θ_h after observing s . Under the true model, the probability that manager’s prediction is correct is given by $p^{m^T}(\theta_h|s)$, but since the manager finds m more decisive than m_T , the manager instead reports a confidence level of $\phi = p^m(\theta_h|s) > p^{m^T}(\theta_h|s)$. Analogously, when s is bad news about productivity, the manager reports an inflated confidence level in her prediction equal to $\phi = p^m(\theta_l|s) > p^{m^T}(\theta_l|s)$. Therefore, when faced with a population of candidates, the manager is systematically overconfident in her predictions about productivity. \blacktriangle

Note that a key condition for overprecision is that $I_{\mathcal{D}}(m_T|s) > I_{\mathcal{D}}(m_{\emptyset}|s)$: under the true model, the realized signal s must (weakly) reduce the DM’s residual uncertainty over the optimal course of action. In other words, the DM is more likely to exhibit overprecision in situations where she is initially uncertain over the optimal course of action. For instance, if we suppose that the hiring manager in the previous example has binary actions $\{\text{hire}, \text{don't hire}\}$, the manager will tend exhibit overprecision when her prior beliefs leave her close to indifferent between hiring and not hiring.

In contrast, consider the case where $I_{\mathcal{D}}(m_T|s) < I_{\mathcal{D}}(m_{\emptyset}|s)$ the signal increases the manager’s residual uncertainty over the optimal course of action. For instance, it may be the case that the manager’s prior beliefs strongly recommend that she hire the candidate, but she receives a signal to the contrary that would push her to greater uncertainty over whether to hire under m_T . In this case, the decisiveness criterion instead predicts that the manager may seek models that *minimize* the informational content of her signal, a form of confirmation bias. The next section studies exactly this implication.

3.2.2 Confirmation Bias

Consider a setting with binary states: $\Theta = \{\theta_l, \theta_h\}$; fix the decision problem \mathcal{D} and signal realization s . Say that m is in favor of θ_h if $m(s|\theta_h) \geq m(s|\theta_l)$, and that m is in favor of θ_l if $m(s|\theta_h) \leq m(s|\theta_l)$. Also, say that a model m is *interior* if $p^m(\cdot|\theta) \neq \delta_{\theta}$ for any state θ . The DM then exhibits confirmation bias in the following sense:

Corollary 2. *For any interior m in favor of θ_h , there exists a threshold prior belief $\bar{p} > 0$ such that if $q(\theta_h) \leq \bar{p}$, $I_{\mathcal{D}}(m'|s) \geq I_{\mathcal{D}}(m|s)$ for any m' in favor of θ_l . Furthermore, for $q(\theta_h) \leq \bar{p}$ and any m'' satisfying $m''(s|\theta_h)/m''(s|\theta_l) \leq m(s|\theta_h)/m(s|\theta_l)$, we also have $I_{\mathcal{D}}(m'|s) \geq I_{\mathcal{D}}(m''|s)$ for any m' in favor of θ_l .*

In words: given model m providing evidence towards one conclusion, there exists a range of prior beliefs favoring the other conclusion under which the DM finds models that agree with her prior more decisive than m . Furthermore, this range of prior beliefs is decreasing in the strength of the evidence under m .

Example (Prior-Based Polarization). Following the classic experimental setup of Darley and Gross (1983), suppose that there are two subjects whose priors disagree about the state $\Theta = \{\theta_h, \theta_l\}$ (in their experiment, the reading ability of a student), but who are otherwise identical. In particular, we have $q_A(\theta_h) > q_B(\theta_h)$: subject A believes θ_h is more likely whereas subject B believes θ_l is more likely.

Suppose that both subjects receive the same signal s (results from an ability test), where there is uncertainty over how to interpret s (e.g. uncertainty over how difficult or diagnostic the test is): both subjects entertain models in favor of θ_h and in favor of θ_l . Assume that the evidence does not completely resolve the debate; all models the DMs entertain are interior.

Corollary 2 states that there exists $\bar{p}, \underline{p} \in (0, 1)$ such that if $q_A(\theta_h) > \bar{p}$ and $q_B(\theta_h) < \underline{p}$, subject A will adopt a model in favor of θ_h and subject B will adopt a model in favor of θ_l — that is, after observing the *same* common signal, beliefs become increasingly polarized. Darley and Gross (1983) and similar studies (Lord et al., 1979; Plous, 1991) document exactly this kind of belief polarization along prior beliefs. \blacktriangle

The decisiveness criterion not only rationalizes the evidence for confirmation bias discussed above, but also produces predictions regarding the when the bias is likely to arise, as the example below illustrates.

Example (Confirmation Bias in Hiring). Consider again the hiring manager example from the discussion of overprecision. Suppose that the hiring manager faces the decision problem

	θ_h	θ_l
<i>hire</i>	v	$-k$
<i>reject</i>	0	0

and that under the true model, $r \equiv m_T(s|\theta_h)/m_T(s|\theta_l) > 1$; under the true model, s is good news about the productivity of the candidate, where r denotes the diagnostic strength of this signal. Note that residual uncertainty is increasing in $p(\theta_h)$ for $p(\theta_h) \in [0, k/(k+v)]$. Bayes' rule implies that so long as $q(\theta_h) < \frac{k}{rv+k}$, $p^{m_T}(\theta_h) < k/(k+v)$, and so for $q(\theta_h) < \frac{rk}{v+rk}$, the hiring manager prefers the uninformative model m_\emptyset over the true model m_T . Intuitively, for $q(\theta_h)$ low enough, the hiring manager is confident that rejecting the candidate is the right call, and so will ignore a disconfirmatory signal that induces greater uncertainty over whether she should hire. \blacktriangle

This analysis demonstrates two key predictions that the decisiveness criterion makes regarding confirmation bias. First, the decisiveness criterion predicts limits to confirmation bias: if under the true model the disconfirmatory signal is sufficiently informative — that is, if r is large — the DM will find the true model more decisive than the uninformative model, and so will not exhibit confirmation bias. In fact, the analysis on overprecision implies that in such cases the DM will instead exhibit overinference — to the extent there are interpretations of the data that overstate the informational content of the signal, the DM will select these interpretations. These limits to confirmation bias depend in an intuitive way on the DM's prior: the more certainty afforded by the DM's prior, the greater scope of disconfirmatory models the DM will find less decisive than the uninformative model.

The second prediction is that the DM's decision problem determines the regions of prior beliefs where confirmation bias operates. This is clearly illustrated in the above example: fixing the signal strength, parameterized by r , the DM exhibits confirmation bias

for priors in the region $q(\theta_h) \in [0, \frac{k}{rv+k}]$. We can see that as v increases — that is, as hiring the candidate becomes a more attractive option — greater certainty in the DM’s prior is required to sustain confirmation bias toward maintaining a negative view of the candidate. Contrast this to approaches to modeling confirmation bias in which the region of priors in which the bias operates is exogenously determined, such as Rabin and Schrag (1999), where confirmation bias occurs for $q \in [0, 0.5)$ and $q \in (0.5, 1]$. Under my model, a disconfirmatory signal is not necessarily one that moves the DM’s beliefs towards 50-50, but rather one that reduces DM’s confidence in the course of action they would have chosen in the absence of the signal.

Taken together, the results on overprecision and confirmation bias imply that the DM’s beliefs will exhibit the following pattern, to the extent the DM entertains models that both overstate and understate the decisiveness of her signal: if the DM is initially uncertain (with respect to the optimal course of action), she will exhibit overreaction to information. However, if the DM’s prior is sufficiently concentrated in one state, she will tend to underreact to disconfirmatory news while continuing to overreact to confirmatory news. As such, the decisiveness criterion predicts the tendency of individuals to both “jump to conclusions” — evaluators form strong impressions based on unreliable diagnostic tools, noise traders act on illusory correlations in price data, and forecasters overstate the accuracy of their estimates — and, once strong convictions have been formed, to explain away disconfirmatory evidence: stereotypes tend to persist even in the face of counterexamples, and initial disagreement can become increasingly polarized in light of new information.

3.3 Comparison to Alternative Selection Criteria

3.3.1 Decisiveness vs. Parsimony

Given the complexity of information decision-makers are often confronted with, it is plausible that individuals may seek parsimonious models of the world to organize data and make decisions, consistent with the view that individuals are “cognitive misers” (Kahneman and Tversky, 1973). Does the decisiveness criterion generate a demand for parsimonious models? Given the difficulty of formalizing a general notion of model complexity, I instead outline key intuitions regarding how the decisiveness criterion relates to a demand for model parsimony through a simple setting, first analyzed in Hong et al. (2007).

The DM is trying to predict whether the price of a stock will increase (θ_h) or decrease (θ_l) in the following period, and observes two sources of news: $s = (s_A, s_B)$, where $s_i \in \{0, 1\}$ for $i = A, B$. Motivated by the view that individuals are “cognitive misers” and so seek simplified models of the world, Hong et al. (2007) analyze a situation in which the DM selects one of two models, each of which uses only a single news source (either A or B) to predict the stock price. In this example, I study when the decisiveness criterion generates demand for such parsimonious models.

Formally, say that model m ignores news source i if for any realization of s_{-i} , posterior beliefs are constant in the value of s_i , i.e.

$$\frac{m(s_i = 1, s_{-i} | \theta_h)}{m(s_i = 0, s_{-i} | \theta_h)} = \frac{m(s_i = 1, s_{-i} | \theta_l)}{m(s_i = 0, s_{-i} | \theta_l)}.$$

First, note that under the unrestricted model space \mathcal{M} , given any signal realization, a model that ignores source i can implement any posterior belief, as can a model utilizing both news sources. As a result, without further structure on the model space, both classes of models are equally privileged under the decisiveness criterion. As such, we impose the following structure on M : suppose that M contains the set of models m satisfying the following conditions:

1. Under m , the s_A, s_B are independent conditional on θ . Let $m(s_i|\theta) = \sum_{s_{-i}} m(s_i, s_{-i}|\theta)$.
2. $m(s_i|\theta) \in [\underline{q}_i, \bar{q}_i] \subset [0, 1]$ for all θ, i .
3. $m(s_i = 1|\theta_h) \geq m(s_i = 1|\theta_l)$ for $i = A, B$.

Condition 2 ensures that that no model completely eliminates residual uncertainty, and Condition 3 implies an ordering on news: $s_i = 1$ implies good news about the stock price, whereas $s_i = 0$ implies bad news. Under these conditions, the signal realization crucially determines whether the DM selects a model that ignores a news source, as the following cases illustrates.

Case 1: Mixed signals. Suppose that $s_A = 1, s_B = 0$ (the analysis is identical for the case where $s_A = 0, s_B = 1$). In this case, the set of models that the DM adopts, for any decision problem, are models that ignore either source A or B . To see this, note that any m that maximizes $p^m(\theta_h)$ must ignore source B , the source delivering bad news; similarly, any m minimizing $p^m(\theta_h)$ must ignore source A , the source delivering good news. Proposition 1 implies that only models implementing these extreme beliefs will be selected under the decisiveness criterion (ignoring ties in the decisiveness ranking) for any decision problem.

Case 2: Aligned signals. Suppose that $s_A = 1, s_B = 1$ (the analysis is identical for the case where $s_A = 0, s_B = 0$). In this case, the set of models the DM adopts, for any decision problem, are models that either ignore *both* sources, or take both sources into account. The logic is similar to the case above: any m that maximizes $p^m(\theta_h)$ must take both sources into account, and any m minimizing $p^m(\theta_h)$ must ignore both sources.

Under the decisiveness criterion, the DM does not value parsimony *per se*, as the contrast between the two cases illustrates: the DM adopts the simplified models considered in Hong et al. (2007) only when the two sources produce conflicting news. Under this account, the DM dislikes incorporating additional dimensions into her model to the extent these dimensions increase her residual decision uncertainty. As such, to the extent that news sources are correlated under the true model and therefore are more likely to agree, the DM will more often adopt models that take these multiple sources into account. However, to the extent news sources are independent and therefore more likely to disagree, the DM will tend to adopt models that selectively attend to certain sources.³

³Note that such models need not be “simple”; a model that assigns differing diagnostic weights to each news source is arguably more complex than a model that treats each source as identical, and yet will tend to be more decisive in the case of mixed signals.

3.3.2 Decisiveness vs. Blackwell Ordering

Note that in my framework, each model m is a Blackwell experiment (Blackwell, 1953), and so can be ranked according to the Blackwell ordering, a measure of the informativeness of an experiment. Formally, m dominates m' in the Blackwell order if for all priors q and decision problems $\mathcal{D} = (A, u)$,

$$\sum_s \sum_\theta \max_{a \in A} u(a, \theta) p^m(\theta|s) p^m(s) \geq \sum_s \sum_\theta \max_{a \in A} u(a, \theta) p^{m'}(\theta|s) p^{m'}(s)$$

where $p^m(s) = \sum_\theta m(s|\theta)q(\theta)$ gives the likelihood of observing s under model m .

While both the decisiveness criteria and the Blackwell order capture notions of the informativeness of a model, note that the decisiveness criteria does not respect the Blackwell order. To see why not, as Section 3.2.2 on confirmation bias demonstrates, a completely uninformative model m_\emptyset may be more decisive than an informative model m , whereas m necessarily dominates m_\emptyset in the Blackwell order. The key distinction driving the discrepancy is that the decisiveness criterion is an ex-post notion of informativeness evaluated for a given signal realization, whereas the Blackwell order is an ex-ante notion. Indeed, it can be shown that if m dominates m' in the Blackwell order, then the *average* decisiveness of m must be greater than that of m' : $\sum_s I_{\mathcal{D}}(m|s)p^m(s) \geq \sum_s I_{\mathcal{D}}(m'|s)p^{m'}(s)$.⁴ As such, even if m is more decisive than m' on average, for some signal realizations it may be the case that m' is more decisive than m .

4 Selection as a Function of Payoffs and Objectives

In this section, I study how model selection under the decisiveness criterion varies with the objectives of the DM. I show how these comparative statics can generate documented context effects and attribution errors, as well as predict novel forms of belief polarization.

4.1 Maximal Payoff Profile Improvements

Consider two decision problems $\mathcal{D} = (A, u)$, $\mathcal{D}' = (A', u)$, where $A' = A \cup \{a'\}$. Say that \mathcal{D}' improves the maximal payoff profile of \mathcal{D} if $u(a', \theta) > \max_{a \in A} u(a, \theta)$ for some θ — that is, \mathcal{D}' is formed from \mathcal{D} by adding an action a' that improves the maximal payoff for at least one outcome. For such payoff profile improvements, let $\Delta u_\theta = \max_{a \in A'} u(a, \theta) - \max_{a \in A} u(a, \theta)$ denote the improvement in the maximal payoff associated with state θ when moving from \mathcal{D} , to \mathcal{D}' .

The following gives a condition on how model selection must respond to such payoff profile improvements, in the case where the payoff profile-improving action a' is not chosen:

Proposition 3. *Fix the menu of models M . Suppose \mathcal{D}' improves the maximal payoff profile of \mathcal{D} via a' . Then, for any $m \in C_{\mathcal{D}}(M|s)$ and $m' \in C_{\mathcal{D}'}(M|s)$, if m' does not recommend a' from \mathcal{D}' , then m', m must satisfy $\sum_\theta p^{m'}(\theta|s)\Delta u_\theta \leq \sum_\theta p^m(\theta|s)\Delta u_\theta$.*

⁴In Section 7.2, I analyze precisely this formulation of ex-ante decisiveness, and contrast its predictions with the ex-post formulation.

To parse this condition, consider the case where a' improves the maximal payoff associated with a single state θ . In this case, the proposition states that $p^{m'}(\theta|s) \leq p^m(\theta|s)$: any change in models must result in lower posterior beliefs in the state for which a' improves the maximal payoff. When a' improves the maximal payoff associated with multiple states, the proposition states that any change in models must result in a lower payoff-weighted average of posterior beliefs in the improved states. For intuition behind this result, recall that the decisiveness criterion favors models that concentrate beliefs in states under which the recommended action will be ex-post optimal. The addition of an unchosen action that increases the maximal achievable payoff for a set of states therefore reduces the decisiveness of models that induce high beliefs in those states.

This comparative static on the DM's choice of model selection has implications for the DM's choice of action: the addition of an unchosen action a leads the decision-maker to select models that recommend against choosing a , as well as similar actions. This implication is most clear in the case of binary states, as summarized in the corollary below:

Corollary 3. *Consider a binary state setting, where $\Theta = \{\theta_h, \theta_l\}$. Suppose \mathcal{D}' improves the maximal payoff associated with θ_h of \mathcal{D} via a' . If a' is not among the actions chosen in \mathcal{D}' , then for any action a chosen from \mathcal{D} and any action a'' chosen from \mathcal{D}' , $u(a, \theta_h) \geq u(a'', \theta_h)$ and $u(a, \theta_l) \leq u(a'', \theta_l)$.*

This prediction can account for findings in the experimental literature in which the addition of an unchosen action reduces subjects' propensity to choose similar actions.

Example (Reason-Based Choice). Consider the following experimental finding from Tversky and Shafir (1992). Subjects decide whether to buy a CD player. In the first treatment (the low-conflict treatment), subjects can choose to either buy a mid-range CD player or to defer the purchase. In the second treatment (the high-conflict treatment), subjects can choose to either buy a mid-range CD player, buy a top-of-the-line CD player, or defer the purchase. They find that the proportion of subjects choosing to defer the purchase is greater in the high-conflict treatment vs. the low-conflict treatment. Tversky and Shafir (1992) interpret these findings as suggesting that in the high-conflict treatment, subjects looking to buy a CD player need to weigh the lower price of the midrange model against the higher quality of the top-of-the-line player, a difficult tradeoff that makes deferring the purchase an easier decision to justify.

To translate this setting into the framework, normalize the DM's payoff from deferring to 0, and suppose that the DM's payoff from purchasing a CD player is $\theta q - k$, where q and k are the quality and price of the player, and θ parameterizes how much the DM weighs quality over price. Suppose that the DM is uncertain about the value of θ , and suppose that this state is binary: $\theta \in \{0, 1\}$. Let (q_l, p_l) and let (q_h, p_h) denote the price and quality of the mid-range and top-of-the-line CD player, respectively, and suppose that $k_h > k_l$ and $v_h > v_l > 0$, where $v_h \equiv q_h - k_h$ and $v_l \equiv q_l - k_l$ denote the net quality of the products.

The decision problems corresponding to the low-conflict and high conflict treatments are

\mathcal{D}_{low} : Low-conflict			\mathcal{D}_{high} : High-conflict		
	$\theta = 1$	$\theta = 0$		$\theta = 1$	$\theta = 0$
<i>midrange</i>	v_l	$-k_l$	<i>top-of-line</i>	v_h	$-k_h$
<i>defer</i>	0	0	<i>midrange</i>	v_l	$-k_l$
			<i>defer</i>	0	0

Suppose further that $k_l/(v_l + k_l) < k_h/(v_h + k_h)$ — that is, conditional on choosing to purchase a CD player, the choice is not obvious — there exists a range of beliefs for which the DM prefers *midrange* over *top-of-line*, and vice versa.

To begin, fix a signal realization s ; we will first show there exists some set of models M such that if the DM entertains M , she will choose *midrange* in the low-conflict treatment and *defer* in the high conflict treatment⁵. First note that by Proposition 1, we can restrict attention to the extreme models \bar{m} , \underline{m} satisfying $\bar{m} \in \arg \max_{m \in M} p^m(\theta = 1|s)$ and $\underline{m} \in \arg \min_{m \in M} p^m(\theta = 1|s)$, respectively; Let $\bar{p} = \max_{m \in M} p^m(\theta = 1|s)$ and $\underline{p} = \min_{m \in M} p^m(\theta = 1|s)$ denote the respective extreme posteriors.

Suppose that M is such that $\underline{p} < k_l/(v_l + k_l)$, $k_l/(v_l + k_l) < \bar{p}$; the DM's choice of model is material for whether she buys a CD player or not. Proposition 3 then implies that if $\bar{p} < k_h/(v_h + k_h)$ — that is, no model recommends *top-of-line* — the addition of *top-of-the-line* must increase the decisiveness of \underline{m} relative to \bar{m} . It can be shown that for some M , this leads the DM to switch from \bar{m} in the low-conflict treatment to \underline{m} in the high-conflict treatment, which in turn induces a switch from *midrange* to *defer*⁶.

Therefore, adding an unchosen alternative, the top-of-the-line player, can cause the DM to switch from the midrange player to deferring the decision via a change in models from \bar{m} to \underline{m} . The addition of the top-of-the-line player makes \bar{m} a less satisfying justification for the DM's decision, as \bar{m} now leaves the DM uncertain about which player to buy, causing the switch to \underline{m} , which makes a comparatively more decisive recommendation that the DM should defer the purchase.

Now, note that if the DM chooses *defer* under in the low-conflict treatment, she must also choose *defer* in the high-conflict treatment. To see this, let m be the model selected under \mathcal{D}_{low} and take any $m' \in M$ that does not recommend *defer* from \mathcal{D}_{high} ; let $p^m = p^m(\theta = 1|s)$, $p^{m'} = p^{m'}(\theta = 1|s)$ denote the respective model-induced posteriors. By assumption, we have $I_{\mathcal{D}_{low}}(m|s) > I_{\mathcal{D}_{low}}(m'|s)$ which implies $p^m v_l < (1 - p^{m'})k_l$; this, along with the maintained assumptions on k and l , implies that $I_{\mathcal{D}_{high}}(m|s) > I_{\mathcal{D}_{high}}(m'|s)$. In line with the intuition behind Proposition 3, the addition of *top-of-line* penalizes models that place higher weight on the improved state $\theta = 1$, which precludes a switch from midrange to defer.

Therefore, for any signal realization, the DM may switch from *midrange* in the low-conflict treatment to *defer* in the high conflict treatment, but must choose *defer* in the

⁵An underlying signal structure in this application could involve, for example, the models specifying different interpretations of an advertisement the DM received regarding the CD players.

⁶The set of \bar{p} and \underline{p} that yield this choice pattern is characterized by the condition $\underline{p} \in \left[\frac{(1-\bar{p})k_l}{v_h}, \frac{(1-\bar{p})k_l + \bar{p}(v_h - v_l)}{v_h} \right]$ which is non-empty given the maintained assumptions $\underline{p} < k_l/(v_l + k_l)$, $k_l/(v_l + k_l) < \bar{p} < k_h/(v_h + k_h)$

high-conflict treatment if she chose *defer* in the low-conflict treatment. As such, fixing any probability distribution governing over signals, the proportion of subjects choosing *defer* in the high-conflict treatment is higher compared to that of the low-conflict treatment, in line with the experimental findings.⁷ ▲

4.2 Reductions in Action Value

Consider two decision problems $\mathcal{D} = (A, u), \mathcal{D}' = (A, u')$. Say that action $a^* \in A$ is uniformly worse in \mathcal{D}' relative to \mathcal{D} if $u'(a^*, \theta) \leq u(a^*, \theta)$ for all θ and $u'(a, \theta) = u(a, \theta)$ for all $\theta, a \neq a^*$. The following gives a condition for how the model must change when an action is made uniformly worse:

Proposition 4. *Suppose a^* is uniformly worse in \mathcal{D}' relative to \mathcal{D} . If m recommends a^* from \mathcal{D}' and m' does not recommend a^* from \mathcal{D} then $m \notin C_{\mathcal{D}}(M|s), m' \in C_{\mathcal{D}}(M|s) \implies m \notin C_{\mathcal{D}'}(M|s)$.*

In words, the decisiveness of any model that recommends a^* must decrease relative to a model that did not recommend a^* when a^* is made uniformly worse. Intuitively, if a^* is made uniformly worse, a model that recommends a^* provides a weaker justification toward its recommended action than does any model recommending a different action.

Proposition 4 implies that the more predisposed the DM is to taking an action, the more they will tend to adopt interpretations of the data that recommend that action, as opposed to interpretations recommending other actions. Note the key distinction between this prediction and that of a model selection criterion based on optimism, in which individuals seek interpretations of the data that increase the perceived value of their chosen course of action: a model can produce a recommendation for an action a by increasing the perceived value of a relative to other actions, without increasing the perceived *absolute* value of a . As such, a DM who is predisposed toward a will find a model decisive if it reduces the attractiveness of competing actions, even if it does not result in a higher perceived value of the chosen action a . This can rationalize patterns that are puzzling from the perspective of a optimism-based criterion, such as individuals adopting interpretations that downplay the credibility of evidence suggesting the safety of a new vaccine, or that point to the ineptitude of a political candidate. Such models do not induce optimism but may induce a sense of certainty over the correct course of action, if for instance the decision-maker is choosing whether to take the vaccine or choosing who to vote for. The following example illustrates the differences between the decisiveness criterion and selection criterion based on optimism.

Example (Wishful Thinking vs. Decisiveness). Consider the following experimental results from Bastardi et al. (2011). They study a group of soon-to-be parents with similar

⁷To formally adapt the model to the analysis of choice probabilities, let $\mu(s)$ denote the objective distribution of signals and let $S_{\mathcal{D}}^a$ denote the set of signals in s under which the DM chooses a in decision problem \mathcal{D} ; to avoid dealing with ties in the definition of choice probabilities, assume that $A_{\mathcal{D}}^m$ and $\arg \max_{m \in M} I_{\mathcal{D}}(m|s)$ are singleton sets. The probability of choosing a in decision problem \mathcal{D} is then $Q(a, \mathcal{D}) \equiv \sum_{s \in S_{\mathcal{D}}^a} \mu(s)$. My results indicate that that $Q(\text{defer}, \mathcal{D}_{\text{low}}) \leq Q(\text{defer}, \mathcal{D}_{\text{high}})$ for any set of models M , where the inequality is strict for some M .

priors — all believe that home care is superior to day care. However, parents face different incentives: some intended to use home care, whereas others intended to use day care. When shown with a study providing evidence in favor of the effectiveness of day care, parents who intended to use home care rated the credibility of the study as low, whereas parents who intended to use day care rated the credibility of the study as high. Can wishful thinking explain these results?

Here, the state is given by $\theta = (\theta^h, \theta^d)$, where $\theta^h = 1$ ($\theta^h = 0$) corresponds to home care being effective (not effective), and likewise for θ^d , which corresponds to the effectiveness of day care. For simplicity, suppose that DMs believe that the quality of day care and home care is independent. In line with the study, suppose that DMs hold common priors over the state. DMs face the menu $A = \{a_h, a_d\}$, a choice between home care and day care. Assume that the payoff from choosing each type of care is greater if that type of care is effective:

$$\begin{aligned} u(a_h, (\theta^h = 1, \theta^d)) &> u(a_h, (\theta^h = 0, \theta^d)) && \text{for all } \theta_d \\ u(a_d, (\theta^h, \theta^d = 1)) &> u(a_d, (\theta^h, \theta^d = 0)) && \text{for all } \theta_h \end{aligned}$$

and that the payoff from choosing a given type of care is independent of whether the alternative is effective or not: $u(a_h, (\theta^h, \theta^d))$ is constant in θ_d and $u(a_d, (\theta^h, \theta^d))$ is constant in θ_h . To rationalize the baseline heterogeneity in choice observed in the experiment, suppose that the DMs who initially choose home care find day care uniformly more costly than DMs who initially choose day care.

Now suppose, as in the experiment, that the DM sees a study s providing evidence that day care is effective. The DM entertains two models; either the study is credible (m) or not (m_\emptyset), where $p^m(\theta^d = 1|s) > p^{m_\emptyset}(\theta^d = 1|s) = q(\theta^d = 1|s)$; assume that the study is uninformative about the effectiveness of home care under both m and m_\emptyset .

First consider the inferences of DMs who chose day care at baseline. Under either model, they would choose day care, but m induces lower regret associated with that choice, and so these DMs adopt m ; they rate the study as credible. Note that a model of wishful thinking would make identical predictions in this case; an agent committed to choosing day care would want to believe that day care is effective.

Now consider the inferences of the DMs who chose home care at baseline. According to Proposition 4, these DMs find m less decisive relative to m_\emptyset than do DMs who choose day care at baseline. These DMs are therefore inclined to adopt m_\emptyset : they rate the study as not credible. Note that wishful thinking *cannot* explain this finding: an agent who simply wishes to be optimistic would have no motive to downplay the credibility of good news regarding an option they don't plan to choose (and would in fact avoid doing so insofar as there are costs to warping one's beliefs, as in Caplin and Leahy (2019)); the decisiveness criterion, however, generates exactly this kind of motive. \blacktriangle

4.2.1 Application: Social Attributions

Consider a set of stylized facts from the psychology literature: when making inferences about others' dispositional traits on the basis of their behavior, individuals tend to ignore the confounding role of situational factors in determining behavior, thus committing the

so-called fundamental attribution error (Ross, 1977). However, there are exceptions to this behavior: when forming inferences about ingroup members on the basis of poor behavior, individuals commit the opposite error – they explain away the behavior as the result of situational factors, as opposed to dispositional traits (Vonk and Konst, 1998).

The decisiveness criterion can rationalize these patterns as a result of differences in the DM’s inferential goals across the two settings.

Example (Social Attributions). The DM is uncertain whether an actor is low-type (θ_l) or high-type (θ_h), with uniform priors. The DM observes the actor’s behavior s , and decides whether to interact with or avoid the actor. The payoffs of each action depend on whether the actor is a stranger or ingroup member, corresponding to the decision problems

	$\mathcal{D}_{str} : \text{Stranger}$			$\mathcal{D}_{ig} : \text{Ingroup}$	
	θ_h	θ_l		θ_h	θ_l
<i>interact</i>	v	$-k$	<i>interact</i>	$v + b$	$-k$
<i>avoid</i>	0	0	<i>avoid</i>	0	0

where $v, k, b > 0$. Here, the DM has a preference for interacting with an ingroup member as opposed to a stranger: *interact* is uniformly improved in \mathcal{D}_{ig} relative to \mathcal{D}_{str} .

The DM entertains the following models: under m , behavior is informative of the actor’s type; under m_\emptyset , behavior is explained away by situational factors, and is thus uninformative about type; and the true model $m_T = \lambda m + (1 - \lambda)m_\emptyset$ reflects uncertainty over whether the actor’s behavior was the result of situational factors or their underlying type.

First consider the DM’s inferences over the stranger. Here, Corollary 1 generates a tendency towards the fundamental attribution error: so long as $I_{\mathcal{D}_{str}}(m_T|s) > I_{\mathcal{D}_{str}}(m_\emptyset|s)$, the DM will select m , thereby attributing the stranger’s behavior entirely to their type and neglecting the confounding role of the situation.

In contrast, consider the DM’s inferences over the ingroup member, in the situation where the ingroup member displayed poor behavior, i.e. $p^m(\theta_h|s) < q(\theta_h)$, and that for this signal realization, m recommends that *don’t interact*, whereas m_\emptyset recommends *interact*. Here, Proposition 4 states that the DM’s stronger preference for interacting with the ingroup member increases the decisiveness of m_\emptyset relative to m when faced with an ingroup member as compared to a stranger; if this difference in preferences b is sufficiently large, the DM will select m_\emptyset , and explain away the behavior with situational factors, thus committing a “reversal” of the fundamental attribution error. Here, the DM’s predisposition toward interacting with the ingroup member increases the decisiveness of models that reinforce this course of action. The decisiveness criterion also predicts an asymmetry in attribution documented in (Vonk and Konst, 1998), where the fundamental attribution error only reverses for inferences over negatively-valenced behavior of ingroup members. In particular, if the ingroup member instead exhibited positive behavior, with $p^m(\theta_h|s) > q(\theta_h)$, the DM will instead find m more decisive. ▲

While the attribution patterns discussed above are consistent with a form of ingroup bias in which the DM seeks to interpret information in a way that is favorable towards the

ingroup, the decisiveness criterion makes additional predictions that cannot be rationalized by such a bias. Under the decisiveness criterion, the DM seeks to explain away poor behavior by ingroups not because she has an innate preference for holding favorable beliefs towards the ingroup, but because this interpretation helps reduce her uncertainty over who to interact with; as such, the DM’s inference will be sensitive to changes to the payoffs from interacting with the ingroup member. For example, consider a decision problem \mathcal{D}'_{st} in which the cost of not interacting with the stranger is higher than in \mathcal{D}_{st} — that is, *don’t interact* is uniformly worse in \mathcal{D}'_{st} relative to \mathcal{D}_{st} .⁸ Proposition 4 predicts that such a shift in payoffs can induce the DM to explain away negative behavior in \mathcal{D}'_{st} just as she does in \mathcal{D}_{ig} , thereby bringing inferences in both decision problems in line with each other.⁹

This analysis can be similarly applied to study discrimination in hiring decisions, and in this setting, predicts a novel channel through which taste-based discrimination can drive belief-based biases against minority applicants. The logic is analogous to the example above — if a hiring manager incurs greater costs to hiring a minority applicant due to discriminatory tastes, they will have a greater inclination against adopting positive interpretations of the data that would recommend a decision to hire — such interpretations lead the manager to greater decision uncertainty. This logic similarly suggests channels for interventions in such a setting — since managers’ inferential biases are not due to the minority status of applicants per se, but rather due to the increased costs of hiring such applicants, the model predicts that providing minority hiring incentives can reverse this particular inferential bias.

4.2.2 Application: Belief Polarization

Proposition 4 highlights a channel for differences in DMs’ objectives to lead to belief polarization.

Example (Payoff-Based Polarization). A society of DMs is uncertain over whether a new disease is dangerous (θ_l) or benign (θ_h) with identical priors. The society contains two Types $i = 1, 2$, and each chooses a level of precaution $\{a_1, a_2, \dots, a_n\}$ to take against the disease, where $u_i(a_k, \theta_h)$ is decreasing in k and $u_i(a_k, \theta_l)$ is increasing in k , i.e. actions are ordered from the lowest to highest level of precaution. Assume that $u_2(a_k, \theta) = u_1(a_k, \theta)$ for all $k < n$, and $u_2(a_n, \theta) = u_1(a_n, \theta) - c$ for $c > 0$: that is, type-2 DMs find it more costly to take the highest precaution level relative to type-1 DMs¹⁰.

Both types of DMs receive a signal s (for instance, news reports that disagree on the severity of the disease, and entertain a set of models M (for instance, each model could

⁸For example, \mathcal{D}_{st} might concern a stranger encountered in another neighborhood, whereas \mathcal{D}'_{st} concerns a stranger who happens to be the DM’s new next-door neighbor.

⁹Note that while a model of ingroup bias may predict that biased beliefs toward outgroups may decrease in severity as the costs of holding such distorted beliefs increase, such a model cannot easily rationalize why such a change in incentives would induce the DM to interpret the behavior of outgroup members in an overly positive light, as in our example.

¹⁰This variation in costs could result from differences in “real” costs — for example, DMs employed in essential jobs may have higher costs to staying in lockdown. This variation could also be due to “social” costs — for example, if taking the vaccine results in the DM being viewed as a non-conformer by her social group.

specify the credibility of each report). Proposition 1 implies that only the models that induce the maximum or minimum belief in θ_l will be adopted; let \bar{m} and \underline{m} denote those models, respectively, and let $p^{\bar{m}}, p^{\underline{m}}$ denote the beliefs those models induce in θ_l .

We are interested in when belief polarization occurs: that is, when one type adopts \bar{m} and the other type adopts \underline{m} . To see how this can arise, notice that there exist interior threshold beliefs \bar{p}, \underline{p} such that when $p(\theta_l) > \bar{p}$ (when $p(\theta_l) < \bar{p}$), all types choose the highest (lowest) level of precaution, and consider the case where \bar{m} and \underline{m} recommend actions at either extreme, i.e. $p^{\bar{m}} > \bar{p}$ and $p^{\underline{m}} < \underline{p}$. Proposition 3 implies that, relative to Type-1 DMs, Type-2 DMs — those that face higher costs of taking high precautions — find \underline{m} relatively more decisive compared to \bar{m} .

In particular, when $p^{\bar{m}} > \bar{p}$ and $p^{\underline{m}} < \underline{p}$, there exists a set of model-induced posteriors $(p^{\underline{m}}, p^{\bar{m}})$ under which belief polarization occurs: Type-1 DMs adopt \bar{m} , and Type-2 DMs adopt \underline{m} . Furthermore, this set is increasing in c with respect to set inclusion.¹¹ That is, the more the society disagrees over their objectives, the greater the scope for polarization. \blacktriangle

Like the decisiveness criterion, a model selection criterion based on optimism can also generate belief polarization through heterogeneity in DMs' objectives, as Caplin and Leahy (2019) demonstrate. Note, however, that wishful thinking cannot generate polarization in the above setting if $u(a, \theta_h) > u(a, \theta_l)$ for all a : that is, if utility is higher when the disease is benign as opposed to dangerous. Here, a DM engaged in wishful thinking has no motive to adopt \underline{m} over \bar{m} — there is no reason for a wishful thinker to discredit news that the disease is likely benign, which is strictly positive news. In many settings, a similar logic appears to hold — for example, it seems difficult to argue that a belief that an impending influx of immigrants will reduce employment opportunities and lower wages, or a belief that the opposing administration will tank economy, represents optimistic thinking. Wishful thinking similarly struggles to explain the apparent effectiveness of mudslinging and negative ad campaigns in generating polarization. The decisiveness criterion, however, can support the adoption of such pessimistic interpretations of the data, so long as they guide the DM's decision-making.

5 Choice Under the Decisiveness Criterion

In this section, I study the implications of model selection under the decisiveness criterion on choice. Fixing a signal realization s : given a utility function u and a set of models M , let $\mathcal{C}(A)$ denote the actions recommended by the decisiveness-maximizing models in M :

$$\mathcal{C}(A) = \bigcup_{m \in \arg \max_{m \in M} I_D(m|s)} \arg \max_{a \in A} \sum_{\theta} u(a, \theta) p^m(\theta|s)$$

in which case we say that \mathcal{C} is *represented* by (u, M) given signal s . I relegate discussion of an axiomatic characterization of \mathcal{C} and its identification properties to Appendix A.5, and

¹¹To see this, let $\bar{v}_l = u(a_n, \theta_l), \bar{v}_h = u(a_n, \theta_h)$ and $\underline{v}_l = u(a_1, \theta_l), \underline{v}_h = u(a_1, \theta_h)$. Under the maintained assumption that $p^{\bar{m}} > \bar{p}$ and $p^{\underline{m}} < \underline{p}$, Type-1 DMs adopt \bar{m} and Type-2 DMs adopt \underline{m} if and only if $(\underline{v}_h - \bar{h})(1 - p^{\bar{m}}) < (\bar{v}_l - \underline{v}_l)p^{\underline{m}}$ and $(\underline{v}_h - \bar{h} + c)(1 - p^{\bar{m}}) > (\bar{v}_l - \underline{v}_l - c)p^{\underline{m}}$. If these conditions hold for c , they must also hold for any $c' > c$.

focus on a key property of \mathcal{C} – an aversion to hedging or diversification, and also study a key comparative static: that if the DM entertains a larger set of models, the DM will be more averse to diversification.

Recall that a key property characterizing model selection under the informativeness criterion is a preference for extreme models, as summarized by Propositions 1 and 2. An immediate consequence of the DM's preference for extreme models is that the DM exhibits a preference for extreme actions, or an aversion to hedging.

Proposition 5. *For $a, a' \in A$, let a'' satisfy $u(a'', \theta) = \lambda u(a, \theta) + (1 - \lambda)u(a', \theta)$ for some $\lambda \in (0, 1)$. If $a'' \in \mathcal{C}(A \cup \{a''\})$, then either $a \in \mathcal{C}(A \cup \{a''\})$ or $a' \in \mathcal{C}(A \cup \{a''\})$.*

Proposition 5 states that the DM exhibits a form of *mixture aversion*: she cannot strictly prefer a hedge between two actions to both of those actions. That is, choice under the decisiveness criterion is averse to a specific form of diversification aversion corresponding to mixtures of acts in utility space. In Appendix A.5, I show that mixture aversion is a key axiom in the behavioral characterization of choice under the decisiveness criterion

In addition to predicting this form of diversification aversion, the model yields the following comparative static: if the DM entertains a larger set of models, the DM will be more averse to a different notion of diversification, which I define below.

For the remainder of this section, restrict to a class of choice problems \mathcal{C} represented by the same utility u . For a menu A , let $\bar{u}^A(\theta) = \max_{a \in A} u(a, \theta)$ for all θ — that is, $\bar{u}^A(\theta)$ gives the maximal payoff in each state among actions in A . We should expect any notion of diversification to stipulate that a diversified action does not deliver payoffs that exceed the maximal payoff profile of the menu, $\bar{u}^A(\theta)$. The following definition is more restrictive, in that it requires that the action delivers payoffs that differ from $\bar{u}^A(\theta)$ by a constant.

Definition (*C-Diversified Action*). Say that a is a *C-diversified action* with respect to A if there exists $k > 0$ such that $u(a, \theta) = \bar{u}^A(\theta) - k$ for all θ . Let H_A collect actions that are *C-diversified* with respect to A .

One example of an action that is *C-diversified* action relative to A is an policy to learn the state at a utility cost k , after which the DM chooses the ex-post optimal action from A . The continuation payoff of this policy is given by $\bar{u}^A(\theta) - k$ in each state θ .

Now consider a comparative notion of diversification aversion.

Definition (*Relative Diversification Aversion*). Say that \mathcal{C}' is *more diversification-averse* than \mathcal{C} if for any menu A , $a \in H_A$, $a \notin \mathcal{C}(A \cup \{a\}) \implies a \notin \mathcal{C}'(A \cup \{a\})$.

That is, \mathcal{C}' is more diversification-averse than \mathcal{C} if \mathcal{C}' never chooses a *C-diversified* action from a menu whenever \mathcal{C} does not. As the proposition below demonstrates, a set inclusion order on the model-implied posteriors the DM entertains characterizes this comparative notion. For a set of models M , let $P_{M|s} = \{p^m(\cdot|s) : m \in M\}$ denote the corresponding set of model-implied posteriors given the signal realization. Say that $P_{M|s}$ is *interior* if no model-implied posterior in $P_{M|s}$ induces certainty in a state.

Proposition 6. *Suppose $\mathcal{C}, \mathcal{C}'$ are represented by $(u, M), (u, M')$, respectively, given signal s . If $co(P_{M|s}) \subseteq co(P_{M'|s})$, then \mathcal{C}' is more diversification-averse than \mathcal{C} . Furthermore, if \mathcal{C}' is more diversification-averse than \mathcal{C} and $P_{M'|s}$ is interior, then $co(P_{M|s}) \subseteq co(P_{M'|s})$.*

Proposition 6 is a natural consequence of Proposition 1, which states that under the decisiveness criterion, the DM selects extreme models. If the set of models the DM entertains expands, she will select yet more extreme models and as a result, choose yet more extreme actions. Proposition 6 also states that the converse is true, so long as the DM does not entertain models that eliminate all residual uncertainty. One immediate consequence of Proposition 6 is a comparative static on the DM's valuation of costly but perfectly revealing information.

Example (Value of Information). Consider a DM who chooses to acquire costly but perfectly revealing information about the state, whose utility is quasilinear and separable in money. Consider the action space $Z \times W$, where Z corresponds to a set of state-contingent prizes and $W \subset \mathbb{R}$ corresponds to money. Suppose that the DM's utility is quasilinear and separable in money: for any action (z, w) , $u((z, w), \theta) = v(z, \theta) + w$.

For any menu A , and cost of information $\kappa < 0$ let a_κ^A denote the action that corresponds to the continuation payoff of acquiring perfectly revealing information about the state at monetary cost κ . Given the quasilinearity assumption, we have $u(a_\kappa^A, \theta) = \bar{u}^A(\theta) - \kappa$: by learning the state prior to choosing from A , the DM achieves the maximal utility achievable in each state, less the cost of the information. Note that a_κ^A delivers precisely the payoffs associated with a \mathcal{C} -diversified action. Proposition 6 then implies that if the DM entertains a larger set of models, the DM will have a lower valuation for the information: formally, for $\mathcal{C}, \mathcal{C}'$ represented by $(u, M), (u, M')$, respectively, if $co(P_M) \subseteq co(P_{M'})$, $a_\kappa^A \notin \mathcal{C}'(A \cup \{a_\kappa^A\}) \implies a_\kappa^A \notin \mathcal{C}(A \cup \{a_\kappa^A\})$, and furthermore if $P_{M'}$ is interior, the converse holds. Intuitively, if the DM entertains a larger set of models, she will tend to hold more extreme beliefs due to Proposition 1; this in turn reduces the value of information. \blacktriangle

Although Proposition 6 only characterizes a notion of relative diversification aversion relating to \mathcal{C} -diversified actions, the result has implications for more general notions of diversification. In particular, consider a less restrictive notion that only stipulates that the action delivers payoffs that do not exceed the maximal payoff profile of a menu:

Definition (Diversified Action). Say that a is a *diversified action* with respect to A if $u(a, \theta) \leq u(a_{max}^A, \theta)$ for all θ . Say that a is a *strictly diversified action* with respect to A if $u(a, \theta) < u(a_{max}^A, \theta)$ for all θ .

A consequence of Proposition 6 is that so long as the set of models the DM entertains includes sufficiently extreme (but interior) models, the DM will never choose a strictly diversified action.

Corollary 4. *If a is strictly diversified with respect to A , then there exists a set of interior models \bar{M} such that for any $M \supseteq \bar{M}$ and \mathcal{C} represented by (u, M) given signal s , $a \notin \mathcal{C}(A \cup \{a\})$.*

Example (Underdiversification). A robust finding is that households tend to hold under-diversified portfolios — that is, they hold fewer securities than are needed to eliminate idiosyncratic risk (see, e.g., Blume and Friend 1975, Kelly 1995, Odean 1999, Vissing-Jorgensen 1999, Polkovnichenko 2005, and Goetzmann and Kumar 2004). A common assumption in modeling under-diversification is that the return distributions of securities are known, and that under-diversification results from non-standard preferences such as cumulative prospect theory (Barberis and Huang 2005) or skewness preferences (Mitton and Vorkink 2007). My framework highlights an additional force that can generate under-diversification: when investors face uncertainty over the return distributions of securities, they favor models that pick out “winners” over models that recommend diversification, as diversification is necessarily likely to be ex-post suboptimal for any realization of returns.

To take a stylized example, consider a DM who chooses how much to allocate between two securities $i = 1, 2$, which deliver monetary payoffs in \mathbb{R} . Suppose the DM has utility over money $v : \mathbb{R} \rightarrow \mathbb{R}$; assume only that v is strictly increasing in money. Let Θ describe the possible joint payoff realizations of two securities, and let $r_{i,\theta}$ denote the monetary payoff that security i delivers in state θ . For $\beta \in [0, 1]$, let a_β denote the action that corresponds to allocating portfolio weight β to security 1; we have $u(a_i, \theta) = v(r_{i,\theta})$ for $i = 1, 2$, and $u(a_\alpha, \theta) = v(\beta r_{1,\theta} + (1 - \beta)r_{2,\theta})$. Here, I impose one assumption on the structure of payoffs: that for each θ , $f_1(\theta) \neq f_2(\theta)$ — that is, with probability 1, the two securities deliver different returns.

Fix any interior β . Note that under these assumptions, a_β is strictly diversified with respect to any menu containing a_0 and a_1 . Corollary 4 then implies that so long the DM entertains a sufficiently extreme set of models, she will choose either a_0 or a_1 over a_β in such a menu.

What characterizes a “sufficiently extreme” model? Here, I give a sufficient condition. Let Θ_1 denote the states where security 1 delivers higher payoffs than security 2, and let Θ_2 denote the states where security 2 delivers higher payoffs. As the proof of Corollary 4 shows, there exists $\epsilon > 0$ such that for any m^* satisfying $p^{m^*}(\Theta_1|s) > 1 - \epsilon$, the DM will not choose the diversified action whenever $m^* \in M$. Intuitively, a model that picks out a security as a winner, such as the m^* defined above, will in general be decisive: such a model necessarily induces a belief that investing in that security will likely be ex-post optimal. On the other hand, models that recommend diversification are in general not decisive, since diversification will necessarily be ex-post suboptimal. ▲

6 Applications

I now discuss two additional applications of the decisiveness criterion. Whereas the preceding analysis took as primitive M , the set of models the DM entertains, in each of these applications I extend the framework by considering two distinct forces that may shape M : the supply of models by expert advisors, and the social exchange of models.

6.1 Certainty in Expert Advice

Evidence suggests that individuals are drawn to advisors who provide more certain advice. As Kahneman (2011) writes,

Experts who acknowledge the full extent of their ignorance may expect to be replaced by more confident competitors, who are better able to gain the trust of clients. An unbiased appreciation of uncertainty is a cornerstone of rationality – but it is not what people and organizations want. (page 263)

Consistent with this view, research in psychology has documented that individuals have a more favorable view toward advisors who make more certain forecasts as opposed to moderate ones. In a stark demonstration, Price and Stone (2004) provide subjects with probabilistic forecasts of a series of events (the likelihood that the price of a stock will increase) made by two advisors, coupled with the realization of each event; the forecasts of one advisor are designed to be more extreme (closer to 0% or 100%) but poorly calibrated compared to those of the other advisor¹². The majority of subjects prefer the extreme advisor, despite the presence of outcome data indicating the superior accuracy of the moderate advisor. Studies using related experimental paradigms corroborate this finding (Yates et al., 1996; Gaertig and Simmons, 2018).

An immediate consequence of the decisiveness criterion is that individuals will be drawn to more certain advice. To see this, consider a setting in which the DM entertains models proposed by a set of advisors. Proposition 1 states that the DM will tend to adopt extreme models; in a binary state case, this amounts to the DM adopting models that place the likelihood of the event close to 0% or 100%. As Kahneman (2011) points out, this tendency for individuals to adopt certain advice can lead to a proliferation of overly certain advice in a competitive market for advisors, even if advisors do not inherently have incentives to provide biased advice. Formally, consider a binary state setting with $\Theta = \{\theta_h, \theta_l\}$. A receiver and a set of senders (advisors) share a common prior over the state and observe a common signal realization s . Each sender i proposes a single model m_i to the receiver, who adopts the most decisive model from $M = \bigcup_i m_i$ given her decision problem \mathcal{D} (in case of a tie, suppose the receiver adopts a random model among the decisiveness-maximizing models in M). Each sender i knows the true model m_T , and their payoffs are given by

$$u_i(m_i) = \begin{cases} v - (p^{m_i} - p^{m_T})^2 & \text{if } m_i \text{ is adopted} \\ 0 & \text{otherwise.} \end{cases}$$

where $p^m \equiv p^m(\theta_h|s)$ denotes the posterior induced by model m . Senders earn a payoff v if their model is adopted, and conditional on this, have incentives to provide accurate advice. First, note that in absence of competition between senders, the sender maximizes utility by proposing the true model m_T , which in turn maximizes the utility of the receiver, who adopts the true model. In what follows, consider the (pure strategy) Nash equilibrium

¹²In particular, the forecasts of the moderate advisor are constructed to be accurately calibrated, and the forecasts of the extreme advisor are constructed by adding 15% to each forecast made by the moderate advisor greater than 50%, and vice versa subtracting 15% from each forecast less than 50%.

of the game in which senders simultaneously propose models. Let M^* denote the set of models that the receiver adopts with positive probability in such any such equilibrium.

Proposition 7. *Suppose that $R_{\mathcal{D}}(p)$ is non-constant in a neighborhood of $p = p^{m_T}$. In any Nash equilibrium with at least two senders, for any $m^* \in M^*$, the receiver's equilibrium posterior belief p^{m^*} must be either $\max\{0, p^{m_T} - \sqrt{v}\}$ or $\min\{1, p^{m_T} + \sqrt{v}\}$.*

Proposition 7 states that at an interior equilibrium posterior belief, the model adopted by the sender must satisfy a “zero-profit” condition – it is the most extreme model the receiver can propose that still yields weakly positive payoffs. Competition causes receivers to ratchet up the extremeness of their proposed models, even though receivers have no inherent incentives to push biased models.¹³

Consider further the special case in which $0 < p^{m_T} - \sqrt{v}, p^{m_T} + \sqrt{v} < 1$, and in which the receiver's decision problem \mathcal{D} is *symmetric*: that is, $R_{\mathcal{D}}(p)$ is symmetric around $p = 1/2$. In this case, it can be shown that if $p^{m_T} \neq 1/2$, the equilibrium posterior belief is unique, and is given by

$$p^{m^*} = \begin{cases} p^{m_T} + \sqrt{v} & p^{m_T} > 1/2 \\ p^{m_T} - \sqrt{v} & p^{m_T} < 1/2 \end{cases}$$

Here, the models the senders propose in equilibrium are an exaggerated version of the truth: if the sender believes that θ_h is more likely than not to occur, she will propose a model that exaggerates the likelihood of θ_h , and vice versa if she believes that θ_l is more likely than not to occur.

6.2 Shared Models and Group Polarization

As discussed in Sections 3.2 and 4.2, the decisiveness criterion can generate belief polarization in response to information — that is, receiving identical information can cause the opinions of individuals to diverge along differences in prior beliefs and/or objectives. In this section, I explore another source of polarization: the exchange of models, or interpretations of the data, within a group.

Consider evidence of so-called *group polarization* — the phenomena in which attitudes following group discussion are more extreme than the prior attitudes held by group members. In one such experiment, Schkade et al. (2000) provide subjects with identical case information and ask subjects to individually rate the severity of punishment appropriate for the defendant. Mock juries are then formed and given the task of deliberating as a group on the punishment severity. The key finding is that group deliberation increases the extremity of punishment ratings: for juries in which individual punishment ratings were high, deliberation tends to increase the group rating relative to the median individual rating within the jury, and vice versa when individual punishment ratings were low. Similar evidence of group polarization has been documented in group judgements along a variety

¹³The condition that $R_{\mathcal{D}}(p)$ is non-constant around p^{m_T} rules out the case in which the receiver finds any model that would deliver weakly positive profits to the sender to be equally decisive.

of other dimensions, such as the appropriate level of risk-taking (Myers and Lamm, 1976) or the degree of racial prejudice (Myers and Bishop, 1976).¹⁴

What can account for these effects of group deliberation? One possibility is that deliberation allows for the aggregation of private information, which can lead to group polarization (Roux and Sobel, 2015). One tension with this explanation, as Roux and Sobel (2015) note, is that subjects are given identical information from which to form their judgements in the standard experimental paradigm used to study group polarization. Glaeser and Sunstein (2009) echo this critique, noting evidence for large shifts in beliefs due to group polarization in settings in which individuals likely have little new knowledge or information to bring to the table, such as debates surrounding climate change or affirmative action. How can group polarization arise in settings where individuals lack private information? Glaeser and Sunstein (2009) analyze a model of “credulous Bayesians” in which individuals overstate the informational content of others’ beliefs — in their model, initial heterogeneity in opinions reflects noise instead of differences in private information, and group polarization arises because individuals misattribute this noise to private information. I propose an alternative account of group polarization, in which initial heterogeneity in opinions does not reflect differences in private information but rather in individuals’ *interpretations* of public information — that is, individuals entertain different models — and that group polarization is driven by the exchange of these interpretations.

The decisiveness criteria provides a simple explanation for why exchanging models can result in more extreme judgments: extreme models tend to be decisive. Here, I adopt the formal framework introduced in Schwartzstein and Sunderam (2022) to study the social exchange of models. Consider a setting with binary states $\Theta = \{\theta_h, \theta_l\}$ (e.g. whether the defendant committed a serious or a mild offense) over which a group of individuals share a common prior q , a common decision problem \mathcal{D} (e.g. deciding the severity of punishment to inflict on the defendant), and observe a public signal s . Suppose that each individual i initially entertains a single model m_i with which to interpret the data, and after group deliberation is exposed to the models entertained by the group, entertaining $M_g \equiv \bigcup_i m_i$.¹⁵ Let p^{m_i} denote individual i ’s posterior belief before group deliberation, where for any model m , $p^m \equiv p^m(\theta_h|s)$.

Consider the effects of deliberation on beliefs. First, note that group deliberation expands the set of models that each individual entertains (from $\{m_i\}$ to M_g); this leads to the adoption of more extreme models (Proposition 1), and in turn, more extreme actions (Proposition 6). Second, under the assumption that the group shares a common prior and decision problem, group deliberation causes the beliefs of individual in the group to converge, as each adopts the same model from M_g after deliberation. That is, exchanging models leads individuals in a group to adopt the extreme interpretations of the data found

¹⁴Importantly, group deliberation has been found to polarize not only group judgements in these settings but also the private judgements of individuals in the group, suggesting that group polarization is not purely driven by distortions caused by the group decision-making process such as a social desirability or the diffusion of responsibility.

¹⁵The assumption that each individual entertains only a single model prior to group discussion is not central to the analysis. Without affecting the analysis, one can instead assume that each individual entertains a set of models M_i and that each individual shares M_i with the group, letting m_i denote the model that individual i finds most decisive from M_i .

within the group.

In this manner, the exchange of models can amplify initial differences in individual judgments between groups, generating group polarization. To see this, suppose that the decision problem \mathcal{D} is symmetric, and let $\underline{p} = \min_{m \in M_G} p^m$, $\bar{p} = \max_{m \in M_G} p^m$ denote the extreme viewpoints held by members of the group prior to deliberation. Let p^* denote the viewpoint held by the group after deliberation. We have

$$p^* = \begin{cases} \bar{p} & \bar{p} + \underline{p} > 1 \\ \underline{p} & \bar{p} + \underline{p} < 1 \end{cases}$$

That is, if the initial viewpoints of a group are skewed in favor of θ_h , the exchange of models within that group will cause the viewpoints of the the group to be increasingly skewed in favor of θ_h , and vice versa if initial viewpoints are skewed toward θ_l .¹⁶

Note that a key prediction of this account of group polarization is that following group deliberation, individuals arrive at beliefs that are not only more extreme, but that also provide greater certainty over the optimal course of action. Importantly, this latter prediction need not hold under an account based on Bayesian information aggregation. Under the Bayesian account, group polarization towards a certain conclusion is purely the consequence of individuals obtaining a greater balance of evidence in favor of that conclusion after deliberation, and in particular occurs regardless of the action implications of that conclusion. Therefore, as demonstrated in Appendix A.2, in certain situations the Bayesian account predicts that group deliberation will push individuals to hold *greater* uncertainty over the optimal course of action. This outlines a key distinction between the predictions of the two accounts: under a Bayesian account, deliberation can cause individuals to become increasingly convinced in the conclusion that the case facts are inconclusive over whether the defendant should be convicted or acquitted, or that there is insufficient evidence to determine whether a climate protection policy should be adopted. Under the present account, on the other hand, group deliberation systematically pushes individuals away from such conclusions.

7 Extensions

As discussed in Section 3, the decisiveness criterion favors extreme models, which induce extreme beliefs. However, it is not the case that decision-makers hold extreme beliefs in every decision context – the models individuals adopt and their resulting conclusions appear to be constrained. In the theory, these constraints are given by M , the set of models the DM is willing to entertain. While M can be identified from choice data, as Appendix A.5, demonstrates, the analyst must otherwise specify M in a given application. This section discusses two approaches to how constraints on the space of models can be endogenized: the first imposes a “cost” to deviating from a default belief, while the second imposes an “entry condition” based on the fit of a default model.

¹⁶Analogous results hold if we relax the assumption that \mathcal{D} is symmetric. In the general case, there exists a $p' \in (0, 1)$ and increasing functions f_1, f_2 , such that $p^* = \bar{p}$ if $f_1(\bar{p} - p') > f_2(p' - \underline{p})$ and $p^* = \underline{p}$ if $f_1(\bar{p} - p') < f_2(p' - \underline{p})$

7.1 Cost of Deviating from Default Beliefs

Consider a constraint on model selection that incorporates a “cost” to selecting models that induce deviations from default beliefs, following Hansen and Sargent (2001) and Caplin and Leahy (2019). Let $p_d \in \Delta(\Theta)$ denote the DM’s default belief over outcomes. Some candidate values of p_d are the DM’s prior, the posterior belief induced by a default model, or the posterior beliefs of a Bayesian observer. The DM adopts a model $m \in \mathcal{M}$ to maximize decisiveness, subject to an additional cost of deviating from the default belief, given by the relative entropy from p_d to model-implied posteriors $p^m(\cdot|s)$:

$$\max_{m \in \mathcal{M}} I_{\mathcal{D}}(m) - \frac{1}{\gamma} \sum_{\theta \in \Theta} p^m(\theta|s) \ln \frac{p^m(\theta|s)}{p_d(\theta)}$$

where γ governs the weight attached to this cost: for $\gamma \rightarrow 0$, the DM adopts the model that implements the default beliefs, whereas as $\gamma \rightarrow \infty$, the DM is unconstrained in her model selection. Appendix A.3 outlines basic properties of this plausibility constraint and draws out its implications in a binary state setting.

7.2 Entry Condition on Models

Here, I develop a constraint on model selection operating through an entry condition on the set of models the DM entertains. Let $m_d \in \mathcal{M}$ denote the DM’s default model. Following Schwartzstein and Sunderam (2021), define the fit of a model given realized signal s as $P(m|s) = \sum_{\theta \in \Theta} m(s|\theta)q(\theta)$. The DM adopts a model that minimizes regret, subject to an entry condition that the chosen model must deliver weakly higher fit than the default model. The DM’s problem is then

$$\min_{m \in M \cup \{m_d\}} \sum_{\theta \in \Theta} (u(a, \theta) - \bar{u}^A(\theta)) p^m(\theta|s) \quad \text{s.t.} \quad P(m|s) \geq P(m_d|s).$$

Proposition 1 in Schwartzstein and Sunderam (2021) shows that the above entry condition is equivalent to the constraint: $P(m|s) \geq P(m_d|s)$ if and only if $p^m(\theta|s) \leq \frac{q(\theta)}{P(m_d|s)} \forall \theta$. That is, the entry condition bounds the extent to which models are allowed to move the DM’s beliefs away from the prior, where the bounds are tighter the better the fit of the default model.

Let $P_{m_d} = \left\{ p \in \Delta(\Theta), p(\theta) \leq \frac{q(\theta)}{P(m_d|s)} \forall \theta \right\}$ denote the set of feasible posteriors given the entry condition, and let $M_{m_d} = \{m \in M : p^m \in P_{m_d}\}$ denote the set of models in $M \cup \{m_d\}$ that survive the entry condition. The DM’s model selection problem can be rewritten as

$$\min_{m \in M_{m_d}} \sum_{\theta \in \Theta} (u(a, \theta) - \bar{u}^A(\theta)) p^m(\theta|s).$$

Note that the results from Sections 3 and 4 have analogs that continue to hold in this extension, once the entry condition on models is accounted for. Appendix A.3 analyzes a binary state example under this plausibility constraint.

8 Discussion

This paper presents a theory of model selection and inference based on the insight that individuals seek decisive models, or models that provide clear decision-making guidance, and studies the implications of this model selection criterion on inference and choice. I conclude by discussing potential extensions and additional applications of the theory.

One limitation of the theory is that it studies model selection as a one-shot procedure. In reality, the set of models decision-makers entertain is often in flux as they are exposed to new models, and decision-makers may revise their working model in light of new information. Section 6 studies one such setting, in which the set of models the decision-maker entertains expands as a result of the social exchange of models. A more complete extension of the theory that considers these dynamics could shed light on how the models individuals adopt change over time, and which models tend to survive the realization of uncertainty. An additional set of applications of the theory is to study its implications for model persuasion (Schwartzstein and Sunderam, 2021), in which senders influence receivers' beliefs by proposing models to interpret known data – under the assumption that receivers select models that they find decisive. While Section 6 analyzes a special case of model persuasion in which senders' preferences are aligned with those of the receiver, a more complete analysis of model persuasion under the decisiveness criterion could shed further light on what models we should expect decision-makers to be exposed to in the presence of persuaders.

References

- Aina, C. (2022). Tailored stories.
- Anscombe, F. J. and Aumann, R. J. (1963). A definition of subjective probability. *Annals of Mathematical Statistics*, 34(1):199–205.
- Bastardi, A., Uhlmann, E. L., and Ross, L. (2011). Wishful thinking: Belief, desire, and the motivated evaluation of scientific evidence. *Psychological Science*, 22(6):731–733.
- Bell, D. E. (1982). Regret in decision making under uncertainty. *Operations Research*, 30(5):961–981.
- Blackwell, D. (1953). Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24(2):265–272.
- Caplin, A. and Leahy, J. V. (2019). Wishful thinking.
- Dana, J., Dawes, R., and Peterson, N. (2013). Belief in the unstructured interview: The persistence of an illusion. *Judgment and Decision-making*, 8(5):512–520.
- Darley, J. M. and Gross, P. H. (1983). A hypothesis-confirming bias in labeling effects. *Journal of Personality and Social Psychology*, 44(1):20–33.
- Dunford, N. and Schwartz, J. T. (1957). *Linear Operators, Part 1: General Theory*. John Wiley Sons.
- Eliasz, K. and Spiegel, R. (2020). A model of competing narratives. *American Economic Review*, 110(12):3786–3816.
- Eyster, E., Li, S., and Ridout, S. (2021). A theory of ex post rationalization.
- Festinger, L. (1957). *A theory of cognitive dissonance*. Stanford university press.
- Frankel, A. and Kamenica, E. (2019). Quantifying information and uncertainty. *American Economic Review*, 109(10):3650–3680.
- Gaertig, C. and Simmons, J. P. (2018). Do people inherently dislike uncertain advice? *Psychological Science*, 29(4):504–520.
- Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153.
- Glaeser, E. L. and Sunstein, C. R. (2009). Extremism and social learning. *Journal of Legal Analysis*, 1(1):263–324.
- Hansen, L. and Sargent, T. J. (2001). Robust control and model uncertainty. *American Economic Review*, 91(2):60–66.

- Harmon-Jones, E., Harmon-Jones, C., and Levy, N. (2015). An action-based model of cognitive-dissonance processes. *Current Directions in Psychological Science*, 24(3):185–189.
- Hayashi, T. (2008). Regret aversion and opportunity dependence. *Journal of Economic Theory*, 139(1):242–268.
- Hong, H., Stein, J. C., and Yu, J. (2007). Simple forecasts and paradigm shifts. *The Journal of Finance*, 62(3):1207–1242.
- Izzo, F., Martin, G. J., and Callander, S. (2021). Ideological competition.
- Kahneman, D. (2011). *Thinking, fast and slow*. Macmillan.
- Kahneman, D. and Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, 80(4):237–251.
- Kruglanski, A. W. and Fishman, S. (2009). The need for cognitive closure. *Handbook of Individual Differences in Social Behavior*, pages 343–353.
- Loomes, G. and Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *The Economic Journal*, 92(368):805–824.
- Lord, C. G., Ross, L., and Lepper, M. R. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*, 37(11):2098–2109.
- Moore, D. A., Tenney, E. R., and Haran, U. (2015). Overprecision in judgment. *The Wiley Blackwell Handbook of Judgment and Decision-making*, 2:182–209.
- Myers, D. G. and Bishop, G. D. (1976). Discussion effects on racial attitudes. *Science*, 169(3947):778–779.
- Myers, D. G. and Lamm, H. (1976). The group polarization phenomenon. *Psychological Bulletin*, 83(4):602–627.
- Plous, S. (1991). Biases in the assimilation of technological breakdowns: Do accidents make us safer? *Journal of Applied Social Psychology*, 21(13):1058–1082.
- Price, P. C. and Stone, E. R. (2004). Intuitive evaluation of likelihood judgment producers: Evidence for a confidence heuristic. *Journal of Behavioral Decision Making*, 17(1):39–57.
- Proulx, T. and Inzlicht, M. (2012). The five “a” s of meaning maintenance: Finding meaning in the theories of sense-making. *Psychological Inquiry*, 23(4):317–335.
- Rabin, M. and Schrag, J. L. (1999). First impressions matter: A model of confirmatory bias. *The Quarterly Journal of Economics*, 114(1):37–82.
- Rockafellar, R. T. (2016). *Convex Analysis*. Princeton University Press.

- Ross, L. (1977). The intuitive psychologist and his shortcomings: Distortions in the attribution process. *Advances in Experimental Social Psychology*, 10:173–220.
- Roux, N. and Sobel, J. (2015). Group polarization in a model of information aggregation. *American Economic Journal: Microeconomics*, 7(4):202–32.
- Sarver, T. (2008). Anticipating regret: Why fewer options may be better. *Econometrica*, 76(2):263–305.
- Schkade, D., Sunstein, C. R., and Kahneman, D. (2000). Deliberating about dollars: The severity shift. *Columbia Law Review*, 100:1139–1175.
- Schwartzstein, J. and Sunderam, A. (2021). Using models to persuade. *American Economic Review*, 111(1):276–323.
- Schwartzstein, J. and Sunderam, A. (2022). Shared models in networks, organizations, and groups.
- Stoye, J. (2011). Axioms for minimax regret choice correspondences. *Journal of Economic Theory*, 146(6):2226–2251.
- Tversky, A. and Shafir, E. (1992). Choice under conflict: The dynamics of deferred decision. *Psychological Science*, 3(6):358–361.
- Vonk, R. and Konst, D. (1998). Intergroup bias and correspondence bias: People engage in situational correction when it suits them. *British Journal of Social Psychology*, 37(3):379–385.
- Yates, J. F., Price, P. C., Lee, J.-W., and Ramirez, J. (1996). Good probabilistic forecasters: The consumer’s perspective. *International Journal of Forecasting*, 12(1):41–56.

APPENDIX

A.1 Model Selection for Prediction Problems

As noted in Section 2, a natural approach to applying the theory in situations where the DM does not face a particular decision problem is to assume that the DM learns from data as if they face a prediction problem. Formally, a *prediction problem* is a decision problem $\mathcal{D} = (A, u)$ where the actions $A = \Delta(\Theta)$ consist of a belief report. For a prediction problem $\mathcal{D} = (\Delta(\theta), u)$, refer to u as a *scoring rule*. Say that a scoring rule u , as well as its associated decision problem \mathcal{D} , is *proper* if for all $p, q \in \Delta(\Theta)$, $\sum_{\theta} u(p, \theta)p(\theta) \geq \sum_{\theta} u(q, \theta)p(\theta)$; that is, a prediction problem is proper if the DM maximizes expected utility by truthfully reporting their beliefs.

Example (Proper Prediction Problems).

Quadratic Loss. Suppose $\Theta = \{0, 1\}$. In this case, we can represent $A = \Delta(\Theta)$ with the probability that $\theta = 1$; consider the scoring rule $u(a, \theta) = (a - \theta)^2$. It is well known that this scoring rule is proper. The decisiveness of a model m in this prediction problem is given by

$$I_{\mathcal{D}}(m|s) = -p^m(1 - p^m)$$

where $p^m \equiv p^m(\theta = 1|s)$. For the prediction problem associated with quadratic loss, the DM adopts the model that minimizes her posterior variance.

Logarithmic Loss. Consider the scoring rule $u(p, \theta) = -\log p(\theta)$. This scoring rule is proper, and the decisiveness of a model m in this prediction problem is

$$I_{\mathcal{D}}(m|s) = -\sum_{\theta} p^m(\theta|s) \ln(p^m(\theta|s))$$

That is, logarithmic loss leads the DM to adopt the model that minimizes posterior entropy.

▲

Here, I characterize model selection under the decisiveness criterion in the case where \mathcal{D} is a proper prediction problem. In particular, I ask whether the restriction to proper prediction problems imposes further properties on model selection beyond the necessary and sufficient properties given in Proposition 2. The answer turns out to be negative, as the result below states.

Proposition 8. *A model choice correspondence C satisfies*

1. *Sen's α, β : If $m \in M \subseteq M'$ and $m \in C(M')$, then $m \in C(M)$. Also, if $m, m' \in C_{\mathcal{D}}(M)$, $M \subseteq M'$ and $m' \in C(M')$ then $m \in C(M')$.*
2. *Continuity: For all $m \in \mathcal{M}$, $\{m' \in \mathcal{M} : m' \in C(\{m, m'\})\}$ and $\{m' \in \mathcal{M} : m \in C(\{m, m'\})\}$ are closed.*

3. *Scale Invariance:* For any $m \in \mathcal{M}$, if m' satisfies $m'(s|\theta) = \lambda m(s|\theta) \forall \theta, \lambda > 0$, then $C(\{m, m'\}) = \{m, m'\}$.
4. *Extremeness:* For $m, m' \in M$, if $m, m' \notin C(M)$, then for any $\lambda \in (0, 1)$, $\lambda m + (1 - \lambda)m' \notin C(M)$.
5. *Certainty Preference:* If $m(s|\theta) = 1$ for any $\theta \in \Theta$ and $m \in M$, then $m \in C(M)$.

if and only if there exists a proper prediction problem \mathcal{D} and signal s such that $C(M) = C_{\mathcal{D}}(M|s)$ for all $M \subseteq \mathcal{M}$.

In other words, the class of proper prediction problems is sufficiently rich so as not to rule out any patterns of model selection that would obtain under the decisiveness criterion.

A.2 Shared Models and Group Polarization: Details

Here, I discuss an example illustrating that Bayesian information aggregation need not result in greater certainty over the state.

Example (Diagnosticity of Evidence). Suppose that a group of N individuals have a common, uniform prior over the state $\Theta = \{\theta_l, \theta_h\}$ (e.g. not guilty vs. guilty) and receive a public signal $s_h \in \{s_h, s_l\}$. Individuals are uncertain over how to interpret s_h , and in particular, entertain two different models that describe the data-generating process: under m_1 , the signal is uninformative over the state, and under m_2 , the signal provides evidence towards θ_h :

$$\begin{aligned} m_1(s_h|\theta_h) &= 1 - m_1(s_h|\theta_l) = q > 1/2 \\ m_2(s_h|\theta_h) &= m_2(s_h|\theta_l) = 1/2 \end{aligned}$$

Suppose that individuals share a common, uniform prior over the two models, and each individual i obtains a private iid signal $\psi_i \in \{1, 2\}$ over the model space, where

$$Pr(\psi_i = 1|m_1) = Pr(\psi_i = 2|m_2) = p > 1/2$$

where the ψ_i are also independent of θ and s . It follows that if a greater number of individuals in the group receive the signal $\psi_i = 2$ than receive $\psi_i = 1$, then information aggregation causes individuals to become increasingly certain that m_2 describes the data-generating process, and therefore increasingly uncertain over the state. \blacktriangle

A.3 Plausibility Constraints on Model Selection: Details

Cost of Deviating: Properties. Denoting $\bar{u}^A(\theta) = \max_{a \in A} u(a, \theta)$ as the maximum payoff achievable for state θ , the DM's problem can be re-written as

$$V = \max_{a \in A, p \in \Delta(\Theta)} \sum_{\theta \in \Theta} (u(a, \theta) - \bar{u}^A(\theta)) p(\theta) - \frac{1}{\gamma} \sum_{\theta \in \Theta} p(\theta) \ln \frac{p(\theta)}{p_d(\theta)}$$

Note that the similarity between this formulation and that of Caplin and Leahy (2019): it is identical to their model of wishful thinking, except the DM’s utility function is replaced by a *normalized* utility function $v(a, \theta) = u(a, \theta) - \bar{u}^A(\theta)$, which normalizes the payoff from taking an action in state θ by the maximal utility attainable in θ . The model is as if the agent seeks to maximize anticipatory normalized utility, given a cost of deviating from the default belief.

Fixing an action a , first order conditions for p imply that an interior solution p must satisfy

$$p(\theta) = \frac{\exp [\gamma(u(a, \theta) - \bar{u}^A(\theta))] p_d(\theta)}{\sum_{\theta' \in \Theta} \exp [\gamma(u(a, \theta') - \bar{u}^A(\theta'))] p_d(\theta')}$$

This expression tells us that the DM’s posterior will be a “tilting” of their default posterior that puts less weight on states associated with higher regret given the chosen action a . Substituting the above expression into DM’s problem yields

$$V = \max_{a \in A} \sum_{\theta \in \Theta} p_d(\theta) \frac{\exp(\gamma u(a, \theta))}{\exp(\gamma \bar{u}^A(\theta))}$$

giving a description of the DM’s choice. To interpret this expression, note that the return to taking an action that pays off in state θ is increasing in $p_d(\theta)$, the DM’s default belief in θ , but also is also decreasing in $\bar{u}^A(\theta)$, the maximal payoff achievable in state θ . The intuition for the latter effect is as follows: all else equal, as the maximal payoff associated with state θ increases, the residual uncertainty associated with that state increases, and so the DM will tend to select models that induce lower posteriors over θ . Note that this comparative static echoes Proposition 3.

Example (Binary State Example: Cost of Deviating). The DM wants to learn whether a fund manager is high-skilled (θ_h) or low-skilled (θ_l), and initially believes either possibility is equally likely: $q(\theta_h) = 1/2$. The DM’s signal is the return of the manager’s fund, which is either high or low: $\{r_h, r_l\}$. The DM’s decision problem is given by

	θ_h	θ_l
hire	v	$-k$
don’t hire	0	0

Suppose that $k > v$, so that given the DM’s prior belief, the DM chooses not to hire the manager. Suppose further that the DM’s default belief is $p_d(\theta_h) = 1/2$, which corresponds to a belief that returns are uninformative about manager skill. Consider the case that a low return is realized. Applying the expression above, the optimal model-implied posterior will be

$$p^m(\theta_h) = \frac{1}{1 + \exp(\gamma v)}$$

Note that this model-implied posterior $p^m(\theta_h)$ is weakly lower than $1/2$, and for $\gamma \rightarrow \infty$, $p^m(\theta_h) \rightarrow 0$, corresponding to the case where the DM faces no costs of deviating from the

default belief, whereas for $\gamma \rightarrow 0$, $p^m(\theta_h) \rightarrow 1/2$, the default belief. \blacktriangle

Example (Binary State Example: Entry Condition on Models). Consider again the fund manager setting. The DM wants to learn whether the fund manager is high-skilled (θ_h) or low-skilled (θ_l), and holds priors $q(\theta_h) = q$. The DM's signal is the return of the manager's fund, which is either high or low: $\{r_h, r_l\}$. The DM's decision problem is again given by

	θ_h	θ_l
hire	v	$-k$
don't hire	0	0

where that $k > v$, so that given the DM's prior belief, the DM chooses not to hire the manager. Suppose that the DM's default model m_d satisfies $m_d(r_l|\theta_l) = m_d(r_l|\theta_h) = \phi$. Suppose again that a low return is realized.

Which models survive the entry condition? Note that only models that satisfy $p^m(\theta_h) \in [1 - (1 - q)/\phi, q/\phi]$ will have greater fit than the default model. In particular, consider the case where $\phi = 1$: here, the default model states that the realized signal of a low return was inevitable, regardless of the skill of the manager. In this case, the only model that survives the entry condition is the default model itself. On the other hand, as $\phi \rightarrow 0$ — that is, as the realized signal becomes increasingly unlikely under the default model — the set of surviving models expands to the entire model space. \blacktriangle

A.4 Ex-Ante Decisiveness

Recall that in the basic framework, model selection occurs ex-post: the DM evaluates each model based on its decisiveness given the signal realization. One might instead imagine an account in which the DM evaluates each model according to an *ex-ante* notion of decisiveness, and adopts a model prior to the signal realization. Here, I provide a formulation of an ex-ante notion of decisiveness, and compare its properties to the ex-post formulation.

As in the basic framework, the DM entertains a set of models M and faces a decision problem \mathcal{D} . Letting $p^m(s) = \sum_{\theta} m(s|\theta)q(\theta)$ denote the likelihood of signal s under model m , let

$$I_{\mathcal{D}}^E(m) = \sum_s I_{\mathcal{D}}(m|s)p^m(s)$$

denote the expected decisiveness of m prior to the signal realization. Consider a DM who, prior to the realization of the signal, selects a model from M that maximizes expected decisiveness; refer to this account of model selection as the *ex-ante* criterion and the account developed in Section 2 as the *ex-post* criterion, and let $C_{\mathcal{D}}^E(M) = \arg \max_{m \in M} I_{\mathcal{D}}^E(m)$ denote the models the DM adopts under the ex-ante criterion.

A.4.1 Properties of Model Selection under Ex-Ante Decisiveness

Extremeness. Recall that a key-property of the ex-post criterion is that it favors extreme models: in particular, as Proposition 1 implies, if $m = \lambda m' + (1 - \lambda)m''$, then for any s ,

$I_{\mathcal{D}}(m|s) \leq \max\{I_{\mathcal{D}}(m'|s), I_{\mathcal{D}}(m''|s)\}$. The ex-ante criterion satisfies an analogous property: if $m = \lambda m' + (1 - \lambda)m''$, then $I_{\mathcal{D}}^E(m) \leq \max\{I_{\mathcal{D}}^E(m'), I_{\mathcal{D}}^E(m'')\}$.

Proposition 9. *For $m, m' \in M$, if $m, m' \notin C_{\mathcal{D}}^E(M)$, then for any $\lambda \in (0, 1)$, $\lambda m + (1 - \lambda)m' \notin C_{\mathcal{D}}^E(M)$.*

The intuition for this extremeness result is similar to that of the ex-ante case: the composite model m'' corresponds to a case where the DM is uncertain over how to interpret the information, which results in lower average decisiveness.

Overprecision and Confirmation Bias. Recall that in the ex-post case, the extremeness property generates a form of overprecision. In particular, letting m_{\emptyset} denote an uninformative model, Corollary 1 states that if $I_{\mathcal{D}}(m_{\emptyset}|s) < I_{\mathcal{D}}(m|s)$, then for any m' satisfying $m = \lambda m' + (1 - \lambda)m_{\emptyset}$, it must be that $I_{\mathcal{D}}(m'|s) \geq I_{\mathcal{D}}(m|s)$. That is, under the ex-post criterion, if the DM finds a model m more decisive than her prior for a given signal realization, she will find a model that overstates the informativeness of m yet more decisive. The ex-ante criterion also generates overprecision, but does so *unconditionally*: for any m' satisfying $m = \lambda m' + (1 - \lambda)m_{\emptyset}$, we have $I_{\mathcal{D}}^E(m') \geq I_{\mathcal{D}}^E(m)$: that is, the DM exhibits overprecision irrespective her prior.

Proposition 10. *For any m satisfying $\lambda m + (1 - \lambda)m_{\emptyset} = m_T$ for some $\lambda \in (0, 1)$, $I_{\mathcal{D}}^E(m) \geq I_{\mathcal{D}}^E(m_T)$.*

In fact, the ex-ante criterion produces a more general form of overprecision: if m' dominates m in the Blackwell order, then $I_{\mathcal{D}}^E(m') \geq I_{\mathcal{D}}^E(m)$ — that is, the ex-ante criterion respects the Blackwell order.¹⁷

An immediate consequence of the fact that overprecision holds unconditionally in the ex-ante case is that, unlike ex-post decisiveness, model selection under ex-ante decisiveness does not generate confirmation bias. Recall that in the ex-post case, Corollary 2 states that if the DM's prior is sufficiently concentrated in one state, the DM will find an uninformative model more decisive than a model providing evidence against that state. In contrast, the ex-ante decisiveness criterion predicts that the DM will never find an uninformative model more decisive than an informative model, regardless of her priors.

A.4.2 Selection as a Function of Objectives under Ex-Ante Decisiveness

Maximal Payoff Improvements. Selection under the ex-post criterion exhibits the following comparative static, formalized in Proposition 3: adding an action to the menu that increases the maximal payoff associated with a set of states but is ultimately not chosen can only result in the DM adopting a model that places lower likelihood on those states. Note that one implication of this property is that the ex-ante model violates IIA — the addition of an unchosen action can induce a change in the model the DM adopts, and in turn lead to a change in the action taken by the DM. The ex-ante model, on the other hand, does not allow for such violations of IIA — the addition of unchosen actions cannot induce a change in

¹⁷It can be shown that if $m = \lambda m' + (1 - \lambda)m_{\emptyset}$, then m is a garbling of m' and so is dominated by m' in the Blackwell order.

model selection. This precludes an analog of Proposition 3 from holding in the ex-ante case.

Reductions in Action Value. As formalized in Proposition 3, in the ex-post case, reducing the payoffs of an action a recommended by model m reduces the decisiveness of m relative to any model that does not recommend a . A similar comparative static holds for the ex-ante model. In particular, suppose model m recommends action a for some signal realization; it must be that reducing the payoffs of a reduces the decisiveness of m relative to any model that does not recommend a for any signal realization.

Formally, consider two decision problems $\mathcal{D} = (A, u)$, $\mathcal{D}' = (A, u')$. Say that action $a^* \in A$ is uniformly worse in \mathcal{D}' relative to \mathcal{D} if $u'(a^*, \theta) \leq u(a^*, \theta)$ for all θ and $u'(a, \theta) = u(a, \theta)$ for all θ , $a \neq a^*$.

Proposition 11. *Suppose a^* is uniformly worse in \mathcal{D}' relative to \mathcal{D} . If $m \in M$ recommends a^* from \mathcal{D}' for some $s \in S$ and $m' \in C_{\mathcal{D}}^E(M)$ does not recommend a^* from \mathcal{D} for any $s \in S$, then $m \notin C_{\mathcal{D}}^E(M) \implies m \notin C_{\mathcal{D}'}^E(M)$.*

In words, both versions of the model share the following feature: the more predisposed the DM is to taking an action, the more they will tend to adopt interpretations of the data that recommend that action, as opposed to interpretations recommending other actions.

Relationship to Optimism. Recall that ex-post decisiveness delivers distinct predictions from a model selection criterion based on optimism: in particular, the ex-post decisiveness of a model is evaluated only based on its implications for the relative, as opposed to absolute, values of actions – a model that induces pessimistic beliefs can still be decisive, so long as it produces a strong recommendation toward one course of action over its alternatives.

This distinction no longer holds in the ex-ante case, however. The ranking over models induced by ex-ante decisiveness is equivalent to a model selection criteria based on ex-ante optimism. To see this, note that the expression for expected decisiveness, for $\mathcal{D} = (A, u)$, can be expressed as

$$I_{\mathcal{D}}^E(m) = \sum_s \left[\max_{a \in A} \sum_{\theta} u(a, \theta) p^m(\theta | s) \right] p^m(s) - \sum_{\theta} \max_{a' \in A} u(a', \theta) q(\theta)$$

As the second term on does not depend on m , selecting the model that maximizes ex-ante decisiveness-maximizing model is equivalent to selecting the model that maximizes expected utility.

A.5 Behavioral Characterization and Identification Results

A key primitive in the framework is the set of models the DM entertains, M . In many situations, however, the set of models the DM entertains is difficult or impossible to directly observe. This raises the following questions: if we are unwilling to make a-priori restrictions on M , does the theory nevertheless make meaningful restrictions on choice, and can M be deduced from choice data?

I take up both of these questions, working in an extended environment that includes objective lotteries over outcomes following Anscombe and Aumann (1963). These results

rely on techniques similar to those in Stoye (2011), which characterizes min-max regret choice correspondences, to provide characterization and identification results for the theory.

A.5.1 Extending the Environment

We extend the DM's decision environment as follows. There is a set of prizes Z . The DM chooses between acts $X \equiv \Delta(Z)^\Theta$, where each act $f \in X$ is a mapping from states to finite objective lotteries over prizes. Let \mathcal{A} denote the collection of finite subsets of X . Take as data the choice correspondence $\mathcal{C} : \mathcal{A} \rightrightarrows X$, which satisfies $\mathcal{C}(A) \subseteq A$ for all $A \in \mathcal{A}$.

Because multiple models can induce the same posterior belief for a given signal realization, it will not in general be possible to identify the set of models the DM entertains; the representation we consider will therefore focus on the set of posteriors P induced by the models the DM entertains. Under the maintained assumptions in Section 2, this set of posteriors will be a closed subset of $\Delta(\Theta)$. Let \mathcal{P} denote the collection of all closed subsets of $\Delta(\Theta)$. Let \mathcal{U} denote the set of utility functions from Z to \mathbb{R} , and extend these to $\Delta(Z)$ by taking expectations¹⁸.

Definition (Decisiveness-Maximizing Representation). A choice correspondence \mathcal{C} has a *decisiveness-maximizing representation* if there exists $(u, P) \in \mathcal{U} \times \mathcal{P}$ such that for all $A \in \mathcal{A}$,

$$\mathcal{C}(A) = \bigcup_{p \in \mathcal{I}(P|A)} \arg \max_{f \in A} \sum_{\theta \in \Theta} u(f(\theta)) p(\theta)$$

where

$$\mathcal{I}(P|A) = \arg \max_{p \in P} \left\{ \max_{f \in A} \sum_{\theta \in \Theta} \left[u(f(\theta)) - \max_{f' \in A} u(f'(\theta)) \right] p(\theta) \right\}.$$

If the above holds, say that (u, P) *represents* \mathcal{C} .

In words, if \mathcal{C} is represented by a utility function and a set of model-implied posteriors, \mathcal{C} chooses the acts that maximizes expected utility with respect to a posterior that maximizes decisiveness. Now, introduce two assumptions that will be needed for identification:

Assumption 1: There exists $z, z' \in Z$ such that $u(z) > u(z')$.

Assumption 2: P does not contain δ_θ for any $\theta \in \Theta$.

Assumption 1 amounts to a non-triviality assumption that ensures that the DM is not indifferent between all acts. Assumption 2 is substantive, and rules out models that induce certainty in a state. Assumptions 1 and 2 are crucial for the identification result that follows, but partial identification of the set of extreme models is still possible when Assumption 2 is relaxed, as I discuss below.

¹⁸In particular, for $u : Z \rightarrow \mathbb{R}$, for each $q \in \Delta(Z)$ let $u(q) = \sum_{z \in Z} u(z)q(z)$.

A.5.2 Relationship to Min-Max Regret Models

This model has a tight relationship with the *min-max regret* model, characterized in Hayashi (2008) and subsequently Stoye (2011), which has the representation

$$\mathcal{C}(A) = \arg \min_{f \in A} \max_{p \in P} \sum_{\theta \in \Theta} \left[\max_{f' \in A} u(f'(\theta)) - u(f(\theta)) \right] p(\theta)$$

Here, the DM chooses to minimize the worst-case expected regret, taken with respect to the set of beliefs P . The decisiveness-maximizing representation is equivalent to a *min-min regret* model, where the DM chooses to minimize the *best-case* expected regret. As such, my characterization result parallels the axiomatization in Stoye, with the appropriate adjustment in axioms to replace the max operator with a min operator. In particular, as I discuss below, while the min-max model is characterized by a preference for hedging, my model is instead characterized by an aversion to hedging.

A.5.3 Characterization Result

With slight of notation, we will let a, b, \dots denote both the objective lottery in $\Delta(Z)$ and the corresponding constant act. For $\theta \in \Theta$, let \mathcal{C}_θ denote the *ex-post* choice function defined by $\mathcal{C}_\theta(A) = \mathcal{C}(\{f(\theta) : f \in A\})$ for all $A \in \mathcal{A}$. Say that f *improves on* A if there exists some θ for which $f = \mathcal{C}_\theta(A \cup \{f\})$, and say that menu A' *improves on* A if there exists some $f \in A'$ that improves on A . Also, for menu A , act f , $\lambda \in [0, 1]$ let $\lambda A + (1 - \lambda)f = \{\lambda f' + (1 - \lambda)f : f' \in A\}$. Consider the following axioms:

Axiom 1 (Independence of Non-Improving Alternatives (INA)). If B does not improve on A , then $\mathcal{C}(A \cup B) \cap A \in \{\mathcal{C}(A), \emptyset\}$.

Axiom 2 (Independence of State-Improving Alternatives (ISA)). For $a, b \in \Delta(Z)$, $\lambda \in (0, 1)$: if A does not improve $\lambda B + (1 - \lambda)a$ nor $\lambda B + (1 - \lambda)b$, and additionally $\mathcal{C}(A \cup \lambda B + (1 - \lambda)a) \cap A \neq \emptyset$ and $\mathcal{C}(A \cup \lambda B + (1 - \lambda)b) \cap A \neq \emptyset$, then $\mathcal{C}(A \cup \lambda B + (1 - \lambda)a) = \mathcal{C}(A \cup \lambda B + (1 - \lambda)b)$.

Axiom 3 (Monotonicity). If $f' \in \mathcal{C}_\theta(f, f')$ for all θ , then $f \in \mathcal{C}(A \cup \{f\}) \implies f' \in \mathcal{C}(A \cup \{f'\})$.

Axiom 4 (Mixture Continuity). For $f, g, h \in X$ and $A \in \mathcal{A}$ with $h \in A$, the sets

$$\begin{aligned} & \{\alpha \in [0, 1] : \alpha f + (1 - \alpha)g \in \mathcal{C}(\{\alpha f + (1 - \alpha)g\} \cup A)\} \\ & \{\alpha \in [0, 1] : h \in \mathcal{C}(\{\alpha f + (1 - \alpha)g\} \cup A)\} \end{aligned}$$

are closed.

Axiom 5 (Mixture Independence). For $\lambda \in (0, 1)$, $g \in X$, $f \in \mathcal{C}(A) \implies \lambda f + (1 - \lambda)g \in \mathcal{C}(\lambda A + (1 - \lambda)g)$.

Axiom 6 (Mixture Aversion). For $f, g \in X, \lambda \in (0, 1)$ s.t. $\{f, g\} \subseteq \mathcal{C}(A)$, $f \in \mathcal{C}(A \cup \{\lambda f + (1 - \lambda)g\})$.

Axiom 7 (Non-Triviality). There exists $z, z' \in Z$ such that $\mathcal{C}(\{z, z'\}) = z$.

Axiom 8 (No Certainty). There exists $f \in X$, $A \in \mathcal{A}$, such that $f \notin \mathcal{C}_\theta(A)$ for all θ , and $f \in \mathcal{C}(A)$.

Axioms 1 and 2 together are a weakening of the standard IIA axioms. This weakening reflects the fact that model selection under the decisiveness criterion, and therefore the evaluation of a given action, depends on the decision problem that the DM faces; Axiom 1 places testable restrictions on the nature of this dependence. In particular, INA states that choice satisfies an IIA property with respect to acts that do not improve the maximal payoff profile — that is, only the addition of acts that improve the maximal payoff profile can induce a change in choice among existing alternatives. On the other hand, ISA states that scaling the payoffs of irrelevant maximal actions cannot induce preference reversals among non-maximal actions.

Example (Violation of INA). Suppose the prize space contains two elements, z_h, z_l , where $\mathcal{C}(\{z_h, z_l\}) = z_h$. Identify each act by the probability it places on prize z_h in each state. Suppose we have the acts f, g, h with

	θ_1	θ_2
f	1	0
g	0	1
h	0	0.9

and suppose that $\mathcal{C}(\{f, g\}) = f$ and $\mathcal{C}(\{f, g, h\}) = g$. Choice in this example exhibits the asymmetric dominance effect — the addition of h , which is dominated by g but not f , causes the DM to switch from f to g . ISA rules out such forms of menu dependence. \blacktriangle

Axiom 3 imposes that choice obeys state-wise dominance. Axiom 4 is a technical condition, and Axiom 7 reflects the assumption that the set of acts is non-trivial, whereas Axiom 8 reflects the assumption that the set of models the DM entertains precludes models that induce complete certainty.

Axiom 5 is an analog of the Independence axiom. It states that choice from a menu is invariant to mixing all acts in the menu with another act. It jointly captures two properties of the theory. The first property concerns choice when the DM is restricted to a single model. In this case, the theory collapses to a case of subjective expected utility (SEU), where the DM evaluates prospects according to a fixed set of subjective beliefs (a fixed model) — in SEU, Independence must hold. The second property is a restriction on how model selection can depend on the DM's decision problem. In particular, since the decisiveness of a model depends only on statewise *differences* between the utility levels of alternatives in the menu, shifting the utility levels that all acts deliver in a state by the same constant cannot induce a change in which model the DM selects. Since in the the-

ory, choice for a fixed model satisfies Independence, and model selection also is invariant to such mixing operations performed on DM's decision problem, choice itself must satisfy Independence.

Example (Violation of Mixture Independence). Suppose again that $Z = \{z_h, z_l\}$ with $\mathcal{C}(\{z_h, z_l\}) = z_h$, and again identify each act by the probability it places on prize z_h in each state. Suppose we have the acts f, g, f', g'

	θ_1	θ_2		θ_1	θ_2
f	0.6	0.4	f'	0.3	0.7
g	0.8	0.2	g'	0.4	0.6

and suppose $\mathcal{C}(\{f, g\}) = g$ and $\mathcal{C}(\{f', g'\}) = f'$. Choice in this example can be rationalized by a model of wishful thinking — the DM places higher beliefs in the state under which she can obtain higher payoffs — θ_1 when the menu is $\{f, g\}$ and θ_2 when the menu is $\{f', g'\}$. Note however, that $f' = \frac{1}{2}f + \frac{1}{2}h$, $g' = \frac{1}{2}g + \frac{1}{2}h$, where h satisfies $h(\theta_1) = z_l$, $h(\theta_2) = z_h$, and so this choice pattern violates Mixture Independence. Note that although that the above mixing operation changes the payoff levels that acts deliver in either state, the payoff differences within states are preserved; it is the latter consideration, rather than the former, that determines the decisiveness of a given model. \blacktriangle .

Axiom 6 imposes that choice satisfies an extremeness property: mixing between two acts cannot result in an act that “improves” on those acts. Axiom 6 directly corresponds to the key extremeness property of the decisiveness criterion formalized in Proposition 1 — the DM's tendency to select extreme models directly translates into a tendency to choose extreme actions – or alternatively a tendency against choosing “diversified” actions or hedging. Note the relationship between Axiom 6 and the Uncertainty Aversion axiom in Gilboa & Schmeidler's (1986) characterization of the Min-max Expected Utility (MEU) model, which embodies the opposite preference towards diversification: if the DM is indifferent between two acts, she must prefer a mixture of those acts to either act. As one might expect, documented choice patterns reflecting ambiguity aversion, which MEU was formalized to rationalize, are inconsistent with Axiom 6, and therefore cannot be rationalized by the theory.

Example (Violation of Mixture Aversion). As before, let $Z = \{z_h, z_l\}$ with $\mathcal{C}(\{z_h, z_l\}) = z_h$, and identify each act by the probability it places on prize z_h in each state. Suppose we have the acts f, g, h

	θ_1	θ_2
f	1	0
g	0	1
h	0.5	0.5

with $\mathcal{C}(\{f, g\}) = \{f, g\}$ but $\mathcal{C}(\{f, g, h\}) = h$. This choice pattern corresponds to a experimental findings for a variant of the classic Ellsberg paradox (Becker and Brownson 1964) in which h corresponds to betting on a black ball being drawn from an urn known to have

50 black and 50 red balls, and f and g correspond to betting on a black and red ball, respectively, being drawn from an urn with an unknown mix of black and red balls. Note that such choice patterns are ruled out by Axiom 6. ▲

Although Axiom 6 is at odds with an established body of evidence for ambiguity-averse preferences in choice settings similar in structure to the example above, there is also a body of evidence for mixture-averse preferences consistent with Axiom 6. For example, Heath and Tversky (1991) demonstrate that aversion to ambiguity reverses when subjects bet in domains in which they have high perceived expertise. In particular, when subjects self-identify as having expertise regarding events for which objective probabilities are unavailable (e.g. outcomes of elections or football matches), they prefer to bet on those events over events with known probabilities, even when the former outcomes are judged by subjects to be equiprobable — behavior which violates Uncertainty Aversion but is consistent with Axiom 6. Viewed through the lens of the theory, subjects who have greater expertise in a given domain may be able to entertain a greater range of arguments (models) for or against any given outcome, which would rationalize the documented mixture-averse behavior. Outside of stylized lab experiments, evidence for behavior consistent with mixture aversion abounds: it is well known that investors tend to hold underdiversified portfolios consisting of too few securities to eliminate idiosyncratic risk (e.g. Mitton and Vorkink 2007) and also tend to exhibit home bias, concentrating their ownership in domestic stocks (French and Poterba 1991). As Section 5 demonstrates, in this model, such behavior can be rationalized as a consequence of uncertainty over the payoff distributions of securities.

The following result states that choice in the model is characterized by Axioms 1–6. Recall that Proposition 1 states that for any decision problem, only extreme models can strictly minimize decisiveness. This implies that in general, only identification of the set of extreme model-induced posteriors from choice data will be possible. The result shows that the set of extreme model-induced posteriors are indeed identified, if we additionally assume Axioms 7 and 8 – that is, if Assumptions 1 and 2 are satisfied.

Theorem 1. *\mathcal{C} satisfies Axioms 1–6 if and only if it has a decisiveness-maximizing representation (u, P) . \mathcal{C} additionally satisfies Axioms 7 and 8 if and only if u and P satisfy Assumptions 1 and 2, respectively, and for any (u', P') representing \mathcal{C} , there exist constants $\alpha > 0$, β such that $u' = \alpha u + \beta$, and $\text{ext}(P') = \text{ext}(P)$.*

The proof of the characterization result in Theorem 1 parallels that of the min-max model established in Stoye (2011), which shows that Axioms 1–5, as well as an Uncertainty Aversion axiom, characterize the min-max regret model. Note, however, that while the extreme set of model-implied posteriors is identified in min-max regret model if Axiom 7 (Non-Triviality) is assumed, the same is not the case in my model. Intuitively, if the DM entertains a model-induced posterior $p \in P$ that places certainty in a state θ – that is, if Assumption 2 does not hold – the DM will never select (outside of cases of indifference) a model that does not place certainty in any state. This precludes the identification of the full set of extreme models in the case where Assumption 2 does not hold, and so the No Certainty axiom, which guarantees that Assumption 2 holds – is required for identification. Below I provide a partial identification result in a setting where Assumption 2 is relaxed.

For a set of model-induced posteriors P , let $ct(P) = \{p \in P : p = \delta_\theta \text{ for some } \theta \in \Theta\}$ denote the set of certainty-inducing posteriors. The following result states that when $ct(P)$ is non-empty it can be identified from choice data, and moreover that choice data identifies whether or not $ct(P)$ is non-empty.

Theorem 2. *\mathcal{C} satisfies Axioms 1–7 if and only if it has a decisiveness maximizing representation (u, P) , where u satisfies Assumption 1. Also, for any (u', P') representing \mathcal{C} , there exists constants $\alpha > 0$, β such that $u = \alpha u' + \beta$, and $ct(P) = ct(P')$.*

Note that in the case where the data identifies that $ct(P) = \emptyset$, Assumption 2 holds and so we can use the identification result in Theorem 1 to identify the full set of extreme models from choice data.

A.5.4 Relative Diversification Aversion

Here, I state a comparative static that corresponds to Proposition 6 in the main text, which states that if the DM entertains a larger set of models, the DM will be more averse to diversification. Consider a family of choice functions $\{\mathcal{C}^k\}_{k \in K}$ that have a decisiveness-maximizing representation, and that share the same utility over prizes. It can be shown that if \mathcal{C} has a decisiveness-maximizing representation, choice over menus of constant acts can be described by a preference relation; let \succeq^* denote the preference relation that describes choice over constant acts for members of $\{\mathcal{C}^k\}_{k \in K}$. For a menu A , let f_{max}^A denote any act that satisfies $f_{max}^A(\theta) \in \{f(\theta) : f \in A, f(\theta) \succeq^* g(\theta) \forall g \in A\}$ — that is, f_{max}^A delivers the most preferred outcome in each state among acts in A .

Definition (*C-Diversified Act*). Say that h is a *C-diversified act* with respect to A if there exists lotteries over prizes $\bar{q}, \underline{q} \in \Delta(Z)$ satisfying $\bar{q} \succeq^* \underline{q}$ such that $\frac{1}{2}f_{max}^A(\theta) + \frac{1}{2}\underline{q} \sim^* \frac{1}{2}h(\theta) + \frac{1}{2}\bar{q}$ for all θ . Let H_A collect acts that are *C-diversified* with respect to A .

As in the definition of a *C*-defined action in the main text, a *C*-diversified act has the following interpretation: to construct a *C*-diversified act, one takes the maximal payoff achievable in each state and reduces that payoff by a constant.

Definition 3 (*Relative Diversification Aversion*). Say that \mathcal{C}' is *more diversification-averse* than \mathcal{C} if for any $A \in \mathcal{A}$, $h \in H_A$, $h \notin \mathcal{C}(A \cup \{h\}) \implies h \notin \mathcal{C}'(A \cup \{h\})$.

That is, \mathcal{C}' is more diversification-averse than \mathcal{C} if \mathcal{C}' never chooses a diversified act from a menu whenever \mathcal{C} does not. I now state the analog of Proposition 6 in the main text.

Theorem 3. *3 Suppose $\mathcal{C}, \mathcal{C}'$ are represented by (u, P) , (u, P') , respectively. If $co(P) \subseteq co(P')$, then \mathcal{C}' is more diversification-averse than \mathcal{C} . Furthermore, if \mathcal{C}' is more diversification-averse than \mathcal{C} and P' satisfies Assumption 2, then $co(P) \subseteq co(P')$.*

A.6 Proofs: Propositions in the Main Text

Proof of Proposition 1. Fix some signal s ; in what follows I drop the dependence on s in the notation, where appropriate. Since $\max_a \sum_y u(a, \theta)p(\theta)$ is convex in p and

$\sum_y \max_{a'} u(a', \theta) p(\theta)$ is linear in p , $R_{\mathcal{D}}$ is concave in p . Bayes' rule implies that for all $\alpha \in (0, 1)$, there exists $\lambda \in (0, 1)$ such that $p^{\lambda m + (1-\lambda)m'} = \alpha p^m + (1-\alpha)p^{m'}$, and so $I_{\mathcal{D}}$ is convex. Therefore, $I_{\mathcal{D}}(\lambda m + (1-\lambda)m') \leq \max\{I_{\mathcal{D}}(m), I_{\mathcal{D}}(m')\}$, and so $m, m' \notin C_{\mathcal{D}}(M) \implies \lambda m + (1-\lambda)m' \notin C_{\mathcal{D}}(M)$, thus establishing the first part of the proposition.

Now additionally suppose that $A^m, A^{m'}$ are disjoint. Toward a contradiction, suppose that $\lambda m + (1-\lambda)m' \in C_{\mathcal{D}}(M)$ for $\lambda \in (0, 1)$. Bayes' rule implies that there exists $\alpha \in (0, 1)$ s.t. $\alpha p^m + (1-\alpha)p^{m'} = p^{\lambda m + (1-\lambda)m'}$. We then have $R_{\mathcal{D}}(p^{\lambda m + (1-\lambda)m'}) \leq \alpha R_{\mathcal{D}}(p^m) + (1-\alpha)R_{\mathcal{D}}(p^{m'})$. Fix $a \in A_{\mathcal{D}}^m$ and $a' \in A_{\mathcal{D}}^{m'}$. The preceding inequality implies that for any $a'' \in A_{\mathcal{D}}^{\lambda m + (1-\lambda)m'}$,

$$\sum_{\theta} u(a'', \theta)(\alpha p^m(\theta) + (1-\alpha)p^{m'}(\theta)) \geq \alpha \sum_{\theta} u(a, \theta)p^m(\theta) + (1-\alpha) \sum_{\theta} u(a, \theta)p^{m'}(\theta)$$

Since $\sum_{\theta} u(a'', \theta)p^m(\theta) \leq \sum_{\theta} u(a, \theta)p^m(\theta)$ and $\sum_{\theta} u(a'', \theta)p^{m'}(\theta) \leq \sum_{\theta} u(a', \theta)p^{m'}(\theta)$, the above implies that $\sum_{\theta} u(a'', \theta)p^m(\theta) = \sum_{\theta} u(a, \theta)p^m(\theta)$ and $\sum_{\theta} u(a'', \theta)p^{m'}(\theta) = \sum_{\theta} u(a', \theta)p^{m'}(\theta)$, which in turn implies that $a'' \in A_{\mathcal{D}}^m, A_{\mathcal{D}}^{m'}$, a contradiction. \square

Proof of Proposition 2. Fix some signal s ; in what follows I drop the dependence on s in the notation, where appropriate. Proposition 1 shows necessity of Extremeness; necessity of Sen's α, β , Certainty Preference, and Scale Invariance are immediate from the definition of the regret criterion.

To see necessity of Continuity, let $E \subset \mathbb{R}^n$ denote a convex, compact superset of $\Delta(\Theta)$. Consider $\tilde{H} : E \rightarrow \mathbb{R}$ defined by $\tilde{H}(p) = \max_{a \in A} \left\{ \sum_{\theta} u(a, \theta)p(\theta) - \sum_p \max_{a' \in A} u(a', \theta)p(\theta) \right\}$. Since \tilde{H} is the maximum over a family of affine functions on a convex, compact set, \tilde{H} is convex. By construction \tilde{H} is bounded from above by 0, and is bounded from below. So \tilde{H} is a proper convex function and is therefore continuous on the relative interior of E ,¹⁹ which in turn implies that H , the restriction of \tilde{H} to $\Delta(\Theta)$, is continuous. Therefore, $R_{\mathcal{D}}(p) = -H(p)$ is continuous, and so $I_{\mathcal{D}}(m) = R_{\mathcal{D}}(p^m)$ is continuous since the map $m \rightarrow p^m$ is continuous.

To show sufficiency, note that Sen's α, β implies that there is a preference relation \succeq where $C_{\mathcal{D}}(M) = \{m \in M : m \succeq m' \forall m' \in M\}$. Since \mathcal{M} is a closed, bounded subset of \mathbb{R}^n , Continuity implies that \succeq can be represented by a continuous utility function $V : \mathcal{M} \rightarrow \mathbb{R}$. Let $\tilde{R} = -V$. Since C satisfies Extremeness, \tilde{R} must be concave. Since C satisfies Certainty Preference, \tilde{R} must attain its minimum value at any m such that $p^m = \delta_{\theta}$ for any $\theta \in \Theta$; normalize this minimum value to 0. Scale Invariance implies that $\tilde{R}(m) = \tilde{R}(m')$ whenever $p^m = p^{m'}$, and so the function $R : \Delta(\Theta) \rightarrow \mathbb{R}$ satisfying $R(p^m) = \tilde{R}(m)$ is well defined.

Bayes' rule implies that for all $\lambda \in (0, 1)$, there exists $\alpha \in (0, 1)$ such that $\lambda p^m + (1-\lambda)p^{m'} = p^{\alpha m + (1-\alpha)m'}$, and so R inherits concavity from \tilde{R} , and for any $\theta \in \Theta$, $R(\delta_{\theta}) = 0$.

¹⁹Theorem 10.4 in Rockafellar, "Convex Analysis"

By Theorem 2 of Frankel and Kamenica (2019), there exists a set of actions A and utility function $u : A \times \Theta \rightarrow \mathbb{R}$ such that $R(p) = \sum_{\theta \in \Theta} p(\theta) \max_a u(a, \theta) - \max_a \sum_{\theta \in \Theta} p(\theta) u(a, \theta)$. \square

Proof of Corollary 1. Follows immediately from Proposition 1. \square

Proof of Corollary 2. Fix some signal s ; in what follows I drop the dependence on s in the notation, where appropriate. In the binary state case, the residual uncertainty can be expressed as $R_{\mathcal{D}}(p)$, where $p = p(\theta_h)$. For any model m , let $p^m \equiv p^m(\theta_h)$. By Proposition 2, we know that $R_{\mathcal{D}}(0) = R_{\mathcal{D}}(1) = 0$, $R_{\mathcal{D}}(p) \geq 0$, and that $R_{\mathcal{D}}$ is continuous. This implies that there exists $p^* > 0$ where $R_{\mathcal{D}}(p)$ is increasing for all $p \leq p^*$. By Bayes' rule, if $m(s|\theta_h)/m(s|\theta_l) < \frac{(1-\bar{p})p^*}{\bar{p}(1-p^*)}$, then $p^m \leq p^*$ so long as $q(\theta_h) \leq \bar{p}$, which in turn implies that $I_{\mathcal{D}}(m') \geq I_{\mathcal{D}}(m)$ for $q(\theta_h) \leq \bar{p}$ and any m' in favor of θ_l since $R_{\mathcal{D}}$ is increasing on $[0, p^*]$.

Finally, note that for $q(\theta_h) \leq \bar{p}$ and any m'' with $m''(s|\theta_h)/m''(s|\theta_l) \leq m(s|\theta_h)/m(s|\theta_l)$, $p^{m''} \leq p^m \leq p^*$ and so $I_{\mathcal{D}}(m') \geq I_{\mathcal{D}}(m'')$. \square

Proof of Proposition 3. Fix some signal s ; in what follows I drop the dependence on s in the notation, where appropriate. Note that $R_{\mathcal{D}'}(p^m) \leq R_{\mathcal{D}}(p^m) + \sum_{\theta} p^m \Delta u_{\theta}$. The inequality results from the fact that the expected regret of m under menu A' must be at most the expected regret associated with continuing to choose the action recommended by m under menu A . Since the actions recommended by m' are in the original menu A , we also have the equality $R_{\mathcal{D}'}(p^{m'}) = R_{\mathcal{D}}(p^{m'}) + \sum_{\theta} p^{m'} \Delta u_{\theta}$. Model m' must deliver lower expected regret than m under \mathcal{D}' , and so we have $R_{\mathcal{D}'}(p^{m'}) \leq R_{\mathcal{D}'}(p^m)$, which in turn yields $\sum_{\theta} p^{m'}(\theta) \Delta u_{\theta} \leq \sum_{\theta} p^m(\theta) \Delta u_{\theta}$. \square

Proof of Corollary 3. Fix some signal s ; in what follows I drop the dependence on s in the notation, where appropriate. Take the actions in A that maximize posterior expected utility for some posterior. These actions can be ordered²⁰ a_1, a_2, \dots, a_K where $u(a_j, \theta_h) > u(a_k, \theta_h)$ and $u(a_j, \theta_l) < u(a_k, \theta_l)$ for $j > k$; the utility-maximizing action is increasing with respect to this order in $p(\theta_h)$. By assumption, no $m' \in C_{\mathcal{D}'}(M)$ recommends a' from \mathcal{D}' , and so Proposition 4 implies that for any $m \in C_{\mathcal{D}}(M)$ and $m' \in C_{\mathcal{D}'}(M)$, $p^{m'}(\theta_h) \leq p^m(\theta_h)$. This in turn implies that any a recommended by m from \mathcal{D} and any a'' recommended by m' from \mathcal{D}' satisfies $u(a, \theta_h) \geq u(a'', \theta_h)$, and $u(a, \theta_l) \leq u(a'', \theta_l)$. \square

Proof of Proposition 4. Fix some signal s ; in what follows I drop the dependence on s in the notation, where appropriate. Note that if m recommends a^* from \mathcal{D}' , then m recommends a^* from \mathcal{D} , since $u'(a^*, \theta) \leq u(a^*, \theta)$ for all θ , and $u'(a^*, \theta) = u(a^*, \theta)$ for all

²⁰This ignores actions that produce identical payoffs, which is without loss of generality.

$\theta, a \neq a^*$. Similarly, if m' recommends $a \neq a^*$ from \mathcal{D} , then m' also recommends $a \neq a^*$ from \mathcal{D}' . Let $k_\theta = u(a^*, \theta) - u'(a^*, \theta)$, and let Θ_+ collect the states for which a^* delivers the maximal payoff. We have

$$\begin{aligned} R_{\mathcal{D}}(p^m) &\leq R_{\mathcal{D}'}(p^m) - \sum_{\theta \in \Theta \setminus \Theta_+} k_\theta p^m(\theta) \\ &\leq R_{\mathcal{D}'}(p^m) \end{aligned}$$

Furthermore, since m' recommends the same action from both $\mathcal{D}, \mathcal{D}'$, we have

$$\begin{aligned} R_{\mathcal{D}'}(p^{m'}) &= R_{\mathcal{D}}(p^{m'}) - \sum_{\theta \in \Theta_+} \min \left\{ k_\theta, u(a^*, \theta) - \max_{a \neq a^*} u(a, \theta) \right\} p^m(\theta) \\ &\leq R_{\mathcal{D}}(p^{m'}) \end{aligned}$$

Since $m \notin C_{\mathcal{D}}(M), m' \in C_{\mathcal{D}}(M)$, $R_{\mathcal{D}}(p^{m'}) < R_{\mathcal{D}}(p^m)$, and so the above two inequalities imply $R_{\mathcal{D}'}(p^{m'}) < R_{\mathcal{D}'}(p^m)$, which in turn implies that $m \notin C_{\mathcal{D}}(M)$. \square

Proof of Proposition 5. Follows directly from Theorem 1. \square

Proof of Proposition 6. Follows directly from Theorem 3. \square

Proof of Corollary 4. Suppose that a is strictly diversified with respect to A . There exists an action \tilde{a} that is C -diversified with respect to A satisfying $u(a, \theta) < u(\tilde{a}, \theta) < \bar{u}^A(\theta)$ for all θ . Let \bar{M} denote any set of interior models for which $\sum_{\theta} u(\tilde{a}, \theta) p^m(\theta|s) < \max_{a \in A} \sum_{\theta} u(a, \theta) p^m(\theta|s)$ for all $m \in \bar{M}$. To see that such a set exists, let $a^* \in A$ be any action that delivers the maximal payoff some state θ^* ; there exists an $\epsilon > 0$ s.t. for any m^* s.t. for any m^* satisfying $p^{m^*}(\theta^*|s) > 1 - \epsilon$, $\sum_{\theta} u(\tilde{a}, \theta) p^{m^*}(\theta^*|s) < \max_{a \in A} \sum_{\theta} u(a, \theta) p^{m^*}(\theta^*|s)$ since $u(\tilde{a}, \theta^*) < \bar{a}^A(\theta^*) = u(a^*, \theta^*)$.

By construction, for \mathcal{C}' represented by (u, \bar{M}) , we have $\tilde{a} \notin \mathcal{C}'(A \cup \{\tilde{a}\})$. By Proposition 6, for any $M \supseteq \bar{M}$ and \mathcal{C} represented by (u, M) , $\tilde{a} \notin \mathcal{C}(A \cup \{\tilde{a}\})$, which implies $a \notin \mathcal{C}(A \cup \{a\})$ as desired. \square

Proof of Proposition 7. Fix any equilibrium, and let M denote the set of models the receiver chooses with positive probability in that equilibrium. To see that for any $m \in M$, we have $p^m \in [\max\{0, p^{m^T} - \sqrt{v}\}, \min\{1, p^{m^T} + \sqrt{v}\}]$, suppose not: any sender who proposed such m incurs negative payoffs, and so can profitably deviate by proposing m' such that $p^{m'} = \max\{0, p^{m^T} - \sqrt{v}\}$, which guarantees a non-negative payoff.

To see that for any $m \in M$, p^m cannot lie strictly in the interior of $[\max\{0, p^{m_T} - \sqrt{v}\}, \min\{1, p^{m_T} + \sqrt{v}\}]$, suppose not. It must then be the case that for all $m \in M$ $p^m \in (\max\{0, p^{m_T} - \sqrt{v}\}, \min\{1, p^{m_T} + \sqrt{v}\})$. To see why, note that by Proposition 2 and the assumption that $R_{\mathcal{D}}(p)$ is non-constant in a neighborhood around p^{m_T} , it must be the case that for any $p \in (\max\{0, p^{m_T} - \sqrt{v}\}, \min\{1, p^{m_T} + \sqrt{v}\})$, $R_{\mathcal{D}}(p) > \min\{R_{\mathcal{D}}(\max\{0, p^{m_T} - \sqrt{v}\}), R_{\mathcal{D}}(\min\{1, p^{m_T} + \sqrt{v}\})\}$. Since the decisiveness of all $m \in M$ must be the same, it cannot be the case that $p^m = \{0, p^{m_T} - \sqrt{v}\}$ or $p^m = \min\{1, p^{m_T} + \sqrt{v}\}$. This in turn implies that senders receive strictly positive profits by proposing a model in M , it must be the case in such an equilibrium that all senders propose a model in M . Let $\underline{p} = \min\{p^{m_T}\} \cup \{p^m : m \in M\}$, $\bar{p} = \max\{p^{m_T}\} \cup \{p^m : m \in M\}$. Consider m', m'' satisfying $p^{m'} = \underline{p} - \epsilon$, $p^{m''} = \bar{p} + \epsilon$, for $\epsilon > 0$; such m', m'' are guaranteed to exist since $p^{m^*} \in (0, 1)$ for all $m^* \in M$. By Proposition 2 and the assumption that $R_{\mathcal{D}}(p)$ is non-constant in a neighborhood around p^{m_T} , the receiver must find at least one of m', m'' strictly more decisive than any model in M^* . Taking $\epsilon \rightarrow 0$, any sender has a profitable deviation to either m' or m'' , a contradiction.

Therefore, for any equilibrium, any model m the receiver chooses with positive probability must satisfy either $p^m = \max\{0, p^{m_T} - \sqrt{v}\}$ or $p^m = \min\{1, p^{m_T} + \sqrt{v}\}$.

Proof of Proposition 9. Note that

$$\begin{aligned} I_{\mathcal{D}}^E(m) &= \sum_s \max_{a \in A} \sum_{\theta} u(a, \theta) p^m(\theta|s) p^m(s) - \sum_{\theta} \max_{a' \in A} u(a', \theta) \rho(\theta) \\ &= \max_{a \in A} \sum_s \sum_{\theta} u(a, \theta) m(s|\theta) \rho(\theta) - \sum_{\theta} \max_{a' \in A} u(a', \theta) \rho(\theta) \end{aligned}$$

and so $I_{\mathcal{D}}^E(m)$ is the maximum over a family of functions that are linear in m . This implies that $I_{\mathcal{D}}^E$ is convex. Therefore, $I_{\mathcal{D}}^E(\lambda m + (1 - \lambda)m') \leq \max\{I_{\mathcal{D}}^E(m), I_{\mathcal{D}}^E(m')\}$, and so $m, m' \notin C_{\mathcal{D}}^E(M) \implies \lambda m + (1 - \lambda)m' \notin C_{\mathcal{D}}^E(M)$ □

Proof of Proposition 10. Follows immediately from the fact that $I_{\mathcal{D}}^E(m)$ respects the Blackwell order over \mathcal{M} , and that for $m = \lambda m_T + (1 - \lambda)m_{\emptyset}$, m_T is a garbling of m and is therefore dominated by m in the Blackwell order. □

Proof of Proposition 11. Note that if m recommends a^* from \mathcal{D}' for some signal s , then m recommends a^* from \mathcal{D} for s , since $u'(a^*, \theta) \leq u(a^*, \theta)$ for all θ , and $u'(a^*, \theta) = u(a^*, \theta)$ for all $\theta, a \neq a^*$. Similarly, if m' recommends $a \neq a^*$ from \mathcal{D} for some signal s , then m' also recommends $a \neq a^*$ from \mathcal{D}' for s . Let S^{a^*} denote the signals for which m recommends a^* .

The above implies that

$$\begin{aligned}
I_{\mathcal{D}'}^E(m') - I_{\mathcal{D}'}^E(m) &= \sum_s \max_{a \in A} \sum_{\theta} u'(a, \theta) p^{m'}(\theta|s) p^{m'}(s) - \sum_s \max_{a \in A} \sum_{\theta} u'(a, \theta) p^m(\theta|s) p^m(s) \\
&= \sum_s \max_{a \in A} \sum_{\theta} u(a, \theta) p^{m'}(\theta|s) p^{m'}(s) - \sum_s \max_{a \in A} \sum_{\theta} u(a, \theta) p^m(\theta|s) p^m(s) \\
&\quad + \sum_{s \in S^{a^*}} [u(a^*, \theta) - u'(a^*, \theta)] p^m(\theta|s) p^m(s) \\
&\geq \sum_s \max_{a \in A} \sum_{\theta} u(a, \theta) p^{m'}(\theta|s) p^{m'}(s) - \sum_s \max_{a \in A} \sum_{\theta} u(a, \theta) p^m(\theta|s) p^m(s) \\
&= I_{\mathcal{D}}^E(m') - I_{\mathcal{D}}^E(m)
\end{aligned}$$

Since $m' \in C_{\mathcal{D}}^E(M)$, $m \notin C_{\mathcal{D}}^E(M)$, $I_{\mathcal{D}}^E(m') - I_{\mathcal{D}}^E(m) \geq 0$, which in turn implies that $I_{\mathcal{D}'}^E(m') - I_{\mathcal{D}'}^E(m) \geq 0$ and $m \notin C_{\mathcal{D}'}^E(M)$. \square

A.7 Proofs: Characterization and Identification Results

Proof of Theorem 1.

Proof. We first prove the characterization result. Necessity of the axioms follow directly from the definition of the representation. The proof of sufficiency proceeds in four steps.

Step 1: Construction of $u : \Delta(Z) \rightarrow \mathbb{R}$.

Define the binary relation \succeq^* on $\Delta(Z)$ as follows: for $a, b \in \Delta(Z)$, $b \succeq^* a$ if $b \in \mathcal{C}(\{b, a\})$. Note that \succeq^* is transitive. To see this, suppose that $c \succeq^* b$, $b \succeq^* a$. We have $b \in \mathcal{C}(\{b, a\})$, $c \in \mathcal{C}(\{c, b\})$. If $a \in \mathcal{C}(\{c, b, a\})$, then by Monotonicity, we have $b \in \mathcal{C}(\{c, b, a\})$ and subsequently $c \in \mathcal{C}(\{c, b, a\})$. On the other hand, if $a \notin \mathcal{C}(\{c, b, a\})$, since a does not improve b , by INA we have $c \in \mathcal{C}(c, b) \implies c \in \mathcal{C}(\{c, b, a\})$. Therefore, we have $c \in \mathcal{C}(\{c, b, a\})$; since b does not improve c , INA implies $c \in \mathcal{C}(\{c, a\}) \implies c \succeq^* a$.

So \succeq^* is a complete and transitive preference relation on $\Delta(Z)$ that agrees with the restriction of \mathcal{C} to menus of constant acts. Furthermore, \succeq^* inherits continuity and independence properties from Mixture Continuity and Mixture Independence: in particular, \succeq^* satisfies

1. For $a, b, c \in \Delta(Z)$, The sets $\{\lambda \in [0, 1] : \lambda a + (1 - \lambda)b \succeq^* c\}$ and $\{\lambda \in [0, 1] : c \succeq^* \lambda a + (1 - \lambda)b\}$ are closed.
2. For $a, b, c \in \Delta(Z)$, $\lambda \in (0, 1)$, $a \succeq^* b \iff \lambda a + (1 - \lambda)c \succeq^* \lambda b + (1 - \lambda)c$.

Therefore, by the Expected Utility Theorem (von-Neumann and Morgenstern 1994) there exists a linear $u : \Delta(Z) \rightarrow \mathbb{R}$ that represents the restriction of \mathcal{C} to menus of constant acts, i.e. for $a, b \in \Delta(Z)$, $u(a) \geq u(b)$ iff $a \in \mathcal{C}(\{a, b\})$.

Step 2: Characterization without Axiom 7.

We will now show that in the case where Non-Triviality does not hold, \mathcal{C} has a decisiveness-maximizing representation. Suppose that Non-Triviality does not hold. This implies that for all $z, z' \in Z$, $u(z) = u(z')$. By linearity of u , we then have that for $a, b \in \Delta(Z)$, $u(a) = u(b)$, and so for any $f, f' \in X$, $f(\theta) \in \mathcal{C}(\{f(\theta), f'(\theta)\})$. Monotonicity then implies that $\mathcal{C}(A) = A$ for all $A \in \mathcal{A}$. Note that by taking any $P \in \mathcal{P}$ we are done; \mathcal{C} has a decisiveness-maximizing representation (u, P) .

The remainder of the proof will deal with the case where Non-Triviality holds. There exists $z, z' \in Z$ s.t. $u(z) > u(z')$; that is, u satisfies Assumption 1. By linearity of u , we can without loss of generality assume that $[-1, 1] \subseteq u(\Delta(Z))$. By INA and Monotonicity, we can identify each act f with its utility-valued act $u \circ f \in u(\Delta(Z))^\Theta$. We will work with utility-valued acts in the remainder of the proof.

Now introduce some notation. Let a_0 denote the constant act delivering 0 utility. For any menu A , let \bar{f}_A denote the act with $\bar{f}_A(\theta) = \max_{f \in A} f(\theta)$, and let $\underline{f}_A(\theta) = \min_{f \in A} f(\theta)$. Let \mathcal{A}_0 denote the set of menus A such that $\bar{f}_A = a_0$, and let X_- denote the set of acts with non-positive utility.

Due to Mixture Independence, we can assume that u is unbounded by taking a homoeothetic extension of \mathcal{C} . Specifically, let $\tilde{\mathcal{A}}$ denote the collection of all finite subsets of \mathbb{R}^Θ , and define the choice function $\tilde{\mathcal{C}}$ on $\tilde{\mathcal{A}}$ by

$$\begin{aligned}\tilde{\mathcal{C}}(A) &= \mathcal{C}(\lambda_A A + (1 - \lambda_A) a_0) \\ \lambda_A &= \max \left\{ 1 / \max_{\theta} \bar{f}_A(\theta), 1 / \max_{\theta} |\underline{f}_A(\theta)| \right\}.\end{aligned}$$

$\tilde{\mathcal{C}}$ is well-defined since $\lambda_A A + (1 - \lambda_A) a_0$ has utility range in $[-1, 1]$ for all $A \in \tilde{\mathcal{A}}$, and extends \mathcal{C} by Mixture Independence.

Step 3: Characterization without Axiom 8.

Suppose No Certainty fails. This implies that for all $b < 0$, $b \notin \mathcal{C}(A)$ for any $A \in \mathcal{A}_0$. Following Step 4 in the proof of Lemma 2 in Stoye (2011), this fact, along with INA, Monotonicity, Mixture Independence, and Mixture Continuity imply that there exists a collection of events $\Sigma_0 \in 2^\Theta$ and $I_0 = \bigcup_{E \in \Sigma_0} \{f \in X_- : \forall \theta \in E, f(\theta) = 0\}$, $\mathcal{C}(A) = A \cap I_0$ for all $A \in X_-$.

Fix any A , $f \in A$, for which there exists $E \in \Sigma_0$ s.t. $f(\theta) = \bar{f}_A(\theta)$ for all $\theta \in E$. We want to show that $f \in \mathcal{C}(A)$. Since $\frac{1}{2}A + \frac{1}{2}(-\bar{f}_M) \in \mathcal{A}_0$, and $\frac{1}{2}f + \frac{1}{2}(-\bar{f}_M) \in I_0$, $\frac{1}{2}f + \frac{1}{2}(-\bar{f}_M) \in \mathcal{C}(\frac{1}{2}A + \frac{1}{2}(-\bar{f}_M))$, which implies $f \in \mathcal{C}(A)$ as desired. By a similar argument, for A , $f \in A$ where there does not exist $E \in \Sigma_0$ s.t. $f(\theta) = \bar{f}_A(\theta)$ for all $\theta \in E$, we have $f \notin \mathcal{C}(A)$.

The above implies that

$$\mathcal{C}(A) = \bigcup_{E \in \Sigma_0} \{f \in A : \forall \theta \in E, f(\theta) = \bar{f}_A(\theta)\}.$$

For each $E \in \Sigma_0$, fix some model induced posterior p_E that has support only on E . Letting $P = \{p_E\}_{E \in \Sigma_0}$, note that (u, P) represents $\mathcal{C}(A)$, since for each $E \in \Sigma_0$ $p_E \in \mathcal{I}(P|A)$ if and only if there exists $f \in A$ for which $f(\theta) \geq g(\theta)$ for all $g \in A$, $\theta \in E$, in which case $f \in \mathcal{C}(A)$.

The remainder of the proof will deal with the case where No Certainty holds.

Step 4: Construction of utility map.

We first state a basic result.

Lemma 1. *There exists $b < 0$, $A \in \mathcal{A}_0$ s.t. $b \in \mathcal{C}(A)$*

Proof. By No Certainty, there exists f, A s.t. $f \notin \mathcal{C}_\theta(A)$ for all θ , and $f \in \mathcal{C}(A)$. We have $\frac{1}{2}A + \frac{1}{2}(-\bar{f}_A) \in \mathcal{A}_0$, and $\frac{1}{2}f + \frac{1}{2}(-\bar{f}_A) < 0$, and by Mixture Independence we have $\frac{1}{2}f + \frac{1}{2}(-\bar{f}_A) \in \mathcal{C}(\frac{1}{2}A + \frac{1}{2}(-\bar{f}_A))$. Fix $b < 0$ such that $b > \frac{1}{2}f + \frac{1}{2}(-\bar{f}_A)$; Monotonicity implies that $b \in \mathcal{C}(\frac{1}{2}A + \frac{1}{2}(-\bar{f}_A) \cup \{b\})$ as desired. \square

Following Stoye (2011), define the binary relation \succeq_C on X_- as follows:

$$\begin{aligned} f \succ_C g &\iff \exists A \in \mathcal{A}_0 \text{ s.t. } f \in \mathcal{C}(A), g \in A \setminus \mathcal{C}(A) \\ f \sim_C g &\iff \exists A \in \mathcal{A}_0 \text{ s.t. } f \in \mathcal{C}(A), g \in \mathcal{C}(A) \end{aligned}$$

Following Step 1 in the proof of Lemma 2 in Stoye (2011), we use INA, Monotonicity, Mixture Continuity, Mixture Independence, and the conclusion of Lemma 1 to collect the following properties of \succeq_C :

1. \succ_C is asymmetric and disjoint from \sim_C and \succeq_C is complete and transitive.
2. \succeq_C is mixture continuous: $\{\lambda : \lambda f + (1 - \lambda)g \succeq_C h\}$ and $\{\lambda : \lambda f + (1 - \lambda)g \preceq_C h\}$ are closed.
3. \succeq_C is homothetic: for all $\lambda < 0$, $f \succeq_C g \iff \lambda f \succeq_C \lambda g$.
4. \succeq_C is monotonic: $f \gg g \implies f \succ_C g$
5. Every act f is \succeq_C -indifferent to exactly one constant act a , which we call the certainty equivalent of f . Let $J : \mathbb{R}^\Theta \rightarrow \mathbb{R}$ map each act $f \in \mathbb{R}^\Theta$ with its certainty equivalent.

Now we collect two additional properties of \succeq_C .

Lemma 2. *\succeq_C is mixture averse: if $f \sim_C g$, then for all $\lambda \in (0, 1)$, $f \succeq_C \lambda f + (1 - \lambda)g$. Furthermore, \succeq_C is C -independent: for $f, g \in X_-$, $b < 0$, $\lambda \in (0, 1)$: $f \succeq_C g \iff \lambda f + (1 - \lambda)b \succeq_C \lambda g + (1 - \lambda)b$.*

Proof. Mixture Aversion of \mathcal{C} straightforwardly implies Mixture Aversion of \succeq_C . To show C -independence, by homotheticity of \succeq_C , it suffices to show that \succeq_C is translation invariant: for all $f, g \in X_-$, $b < 0$, $f \succeq_C g \iff f + b \succeq_C g + b$.

Fix any $f \in X_-$, $b < 0$, and let a be the (unique) certainty equivalent of f . By Lemma 1, there exists $\underline{b} < 0$, $A \in \mathcal{A}_0$ such that $\underline{b} \in \mathcal{C}(A)$; by Mixture Independence, we can without loss assume that $\underline{b} < \min_{\theta} \underline{f}_A(\theta) + b$ and $\underline{b} < a + b$. Monotonicity implies that

$$\begin{aligned} f &\in \mathcal{C}(A \cup \{f\}) \\ a &\in \mathcal{C}(A \cup \{a\}) \\ f + b &\in \mathcal{C}(A \cup \{f + b\}) \\ a + b &\in \mathcal{C}(A \cup \{a + b\}). \end{aligned}$$

In conjunction with INA and disjointedness of \sim_C and \succ_C , the first two lines imply that $\{f, a\} \subseteq \mathcal{C}(A \cup \{f, a\})$; Mixture Independence then implies $\{f + b, a + b\} \subseteq \mathcal{C}(A + b \cup \{f + b, a + b\})$. In conjunction with INA, the last two lines imply that either $\{f + b, a + b\} \cap \mathcal{C}(A \cup \{f + b, a + b\}) \neq \emptyset$. ISA then implies that $C(A + b \cup \{f + b, a + b\}) \cap \{f + b, a + b\} = C(A \cup \{f + b, a + b\}) \cap \{f + b, a + b\}$, and so $f + b \sim_C a + b$.

The above implies that \succeq_C is translation invariant: fixing $f, g \in X_-$, $b < 0$, we have $f \succeq_C g \iff a_f \succeq_C a_g \iff a_f + b \succeq_C a_g + b \iff f + b \succeq_C g + b$. \square

Note that

$$C(A) = \arg \max_{f \in A} J(f - \bar{f}_A).$$

To see this, suppose $f \in \mathcal{C}(A)$. By Mixture Independence, $f - \bar{f}_A \in C(A - \bar{f}_A)$, which in turn implies $f - \bar{f}_A \succeq_C g - \bar{f}_A$ for all $g \in A$, and so $J(f - \bar{f}_A) \geq J(g - \bar{f}_A)$ for all $g \in A$. Conversely, suppose $J(f - \bar{f}_A) < J(g - \bar{f}_A)$ for some $g \in A$. Then $f - \bar{f}_A \prec_C g - \bar{f}_A$, and so by disjointedness of \succ_C and \sim_C , we have $f - \bar{f}_A \notin \mathcal{C}(A - \bar{f}_A)$ and so $f \notin \mathcal{C}(A)$ by Mixture Independence.

Since \succeq_C is mixture continuous, monotonic, C -independent, and mixture averse, following arguments in Gilboa and Schmeidler (1989), there exists a set of posteriors $P \in \mathcal{P}$ such that

$$J(f) = \max_{p \in P} \sum_{\theta} f(\theta) \cdot p(\theta)$$

and so by the above, we have

$$C(A) = \arg \max_{f \in A} \max_{p \in P} \sum_{\theta} \left[u(f(\theta)) - \max_{g \in A} u(g(\theta)) \right] p(\theta).$$

and so (u, P) represents \mathcal{C} . Furthermore, P satisfies Assumption 2, since otherwise \mathcal{C} violates No Certainty.

We now prove the identification result. Suppose that (u, P) and (u', P') satisfy Assumptions 1 and 2 and represent \mathcal{C} . Since the restriction of \mathcal{C} to menus constant acts has an

expected-utility representation, utility is identified up to affine transformation, i.e. there exists $\alpha > 0$, β s.t. $u = \alpha u' + \beta$. We can therefore without loss assume that $u' = u$. Say that $P \in \mathcal{P}$ represents \mathcal{C} if (u, P) represents \mathcal{C} . Since u satisfies Assumption 1, its range is a non-trivial interval $I \equiv u(\Delta(Z))$, and we can without loss assume that $[-1, 1] \subseteq I$. Henceforth, identify each act f by the corresponding utility act $u \circ f \in I^\Theta$.

For $P \in \mathcal{P}$ and c in the strict interior of I , define

$$B(c|P) = \bigcap_{p \in P} \{f \in I^\Theta : f \cdot p < c\}$$

$$\overline{B}(c|P) = \bigcap_{p \in P} \{f \in I^\Theta : f \cdot p \leq c\}$$

Lemma 3. *If P and P' represent \mathcal{C} , then $\overline{B}(c|P) = \overline{B}(c|P')$*

Proof. Fix any $c \in I$, and any $P \in \mathcal{P}$ that represents \mathcal{C} . For $\lambda \in (0, 1)$, $f \in I^\Theta$, let $f_\lambda = (1 - \lambda)f + \lambda c$. We will show that $f \notin B(c|P)$ iff there exists $\lambda \in [0, 1)$ and $A \in \mathcal{A}$ containing c, f_λ such that $f_\lambda \in \mathcal{C}(A)$. To see (\implies), suppose $f \in \mathcal{C}(\{f, c\})$. Then there must exist a model induced posterior $p \in P$ s.t. $f \cdot p \geq c$, and so $f \notin B(c|P)$.

To see (\impliedby), suppose $f \notin B(c|P)$. There exists $p^* \in P$ s.t. $f \cdot p^* \geq c$. Take any sequence $\epsilon_k \rightarrow 0$, where $\epsilon_k \in (0, 1 - c)$ for all k . For all k , there exists $\lambda_k \in (0, 1)$ satisfying $\max_{\theta \in \Theta} f_{\lambda_k}(\theta) < c + \epsilon_k$; fix such a sequence of λ_k . For each k , define a set of acts $G_k = (g_{k,\theta})_{\theta \in \Theta}$ where for each θ, k ,

$$g_{k,\theta}(\theta') = \begin{cases} 0 & \theta \neq \theta' \\ c + \epsilon_k & \theta = \theta' \end{cases}$$

Define a sequence of menus $A_k = G_k \cup \{f_{\lambda_k}, c\}$. By construction, for all A_k we have $\mathcal{I}(P|A_k) = \arg \max_{p \in P} \max_{f \in A_k} f \cdot p$. Since P is closed and no $p \in P$ places probability 1 in any state by Assumption 2, there exists some K such that for all $k > K$, $\max_{p \in P} g_{k,\theta} \cdot p \leq c$ for all θ , which in turn implies that for all $k > K$,

$$\max_{p \in P} g_{k,\theta} \cdot p \leq c \leq f_{\lambda_k} \cdot p^*$$

which in turn implies that for all $k > K$, we have $f_{\lambda_k} \in \arg \max_{f \in A_k} f \cdot p$, and so $f_{\lambda_k} \in \mathcal{C}(A_k)$ as desired.

Now, suppose (u, P) and (u, P') represent \mathcal{C} and fix any $c \in I$. The above implies that $B(c|P) = B(c|P')$, and so $\overline{B}(c|P) = cl(B(c|P)) = cl(B(c|P')) = \overline{B}(c|P')$. \square

Let $\mathcal{P}_{ext} = \{ext(P) : P \in \mathcal{P}\}$ denote the collection of extreme points formed from elements in \mathcal{P} , and let \mathcal{P}_{conv} collect the convex elements in \mathcal{P} . First, we will show that $co : \mathcal{P}_{ext} \rightarrow \mathcal{P}_{conv}$ is one-to-one. To see this, note that for $P \in \mathcal{P}_{ext}$, we have $ext(co(P)) = P$. Also, for $P \in \mathcal{P}_{conv}$, since P is a convex, closed subset of $\Delta(\Theta)$ and is therefore compact, by the Minkowski-Caratheodory Theorem $co(ext(P)) = P$. Therefore, $co : \mathcal{P}_{ext} \rightarrow \mathcal{P}_{conv}$ has an inverse and so is one-to-one.

Now, towards a contradiction, suppose that $ext(P) \neq ext(P')$. Define $P_{conv} = co(ext(P))$ and $P'_{conv} = co(ext(P'))$; by the above, we have $P_{conv} \neq P'_{conv}$. Since P, P' are closed and

therefore compact, so are P_{conv}, P'_{conv} . Since $P_{conv} \neq P'_{conv}$, we can without loss of generality take there to be $p \in P_{conv} \setminus P'_{conv}$. By a separating hyperplane theorem (Dunford and Schwartz, 1957, Theorem V.2.10), there exists a nonzero $\xi \in \mathbb{R}^\Theta$ and a $c \in \mathbb{R}$ such that

$$\max_{p' \in P'_{conv}} \xi \cdot p' < c < \xi \cdot p$$

Without loss of generality, we can take an affine transformation of ξ and c so that $\xi \in I^\Theta$ and $c \in I$. This then implies that $B(c|P_{conv}) \neq B(c|P'_{conv})$. By Proposition 1, the fact that (u, P) and (u, P') represent \mathcal{C} implies that (u, P_{conv}) and (u, P'_{conv}) also represent \mathcal{C} . Lemma 3 implies $B(c|P_{conv}) = B(c|P'_{conv})$, a contradiction. \square

Proof of Theorem 2.

Proof. The proof for the characterization result is contained in Steps 1–2 from the proof for Theorem 1.

To show the identification result, fix some $(u, P), (u', P')$ that represent \mathcal{C} . Following the proof of Theorem 1, we can without loss assume that $u' = u$. Say that $P \in \mathcal{P}$ represents \mathcal{C} if (u, P) represents \mathcal{C} . Since u satisfies Assumption 1, its range is a non-trivial interval $I \equiv u(\Delta(Z))$, and we can without loss assume that $[-1, 1] \subseteq I$. Henceforth, identify each act f by the corresponding utility act $u \circ f \in I^\Theta$.

Now fix any $c \in I$, $\theta \in \Theta$. Take any sequence $\epsilon_k \rightarrow 0$ satisfying $\epsilon_k \in (0, 1 - c)$ for all k . Define a sequence of utility acts $g_{k,\theta}$ where for each k ,

$$g_{k,\theta}(\theta') = \begin{cases} 0 & \theta \neq \theta' \\ c + \epsilon_k & \theta = \theta' \end{cases}$$

and define $A_k = \{c, g_{k,\theta}\}$. It suffices to show that for any P representing \mathcal{C} , $\delta_\theta \notin P$ if and only if there exists K s.t. for all $k > K$, $g_{k,\theta} \notin \mathcal{C}(A_k)$. To see this, suppose that $\delta_\theta \notin P$. Since P is closed, there exists some $\bar{p} < 1$ such that $\max_{p \in P} p(\theta) \leq \bar{p} < 1$ and so for all $p \in P$, $g_{k,\theta} \cdot p \leq \bar{p}(c + \epsilon_k)$. Since $\bar{p} < 1$ and $\epsilon_k \rightarrow 0$, there exists K s.t. for all $k > K$, $\bar{p}(c + \epsilon_k) < c$, which implies that $g_{k,\theta} \notin \mathcal{C}(A_k)$. Conversely, suppose that $\delta_\theta \in P$. This implies $\delta_\theta \in \mathcal{I}(P|A)$ for any menu A , and so $g_{k,\theta} \in \mathcal{C}(A_k)$ for all k . \square

Proof of Theorem 3.

Proof. Begin by proving the first statement. Suppose $(u, P), (u, P')$ represent $\mathcal{C}, \mathcal{C}'$, respectively, and suppose $co(P) \subseteq co(P')$. Fix any $A \in \mathcal{A}$, $h \in H_A$, and suppose $h \notin \mathcal{C}(A, \{h\})$. By construction, there exists $c \leq 0$ such that $u(h(\theta)) - \max_{f \in A} u(f(\theta)) = c$ for all θ . $h \notin \mathcal{C}(A, \{h\})$ implies that

$$\begin{aligned} c &< \max_{p \in P} \max_{g \in A} \left[u(g(\theta)) - \max_{f \in A} u(f(\theta)) \right] \\ &\leq \max_{p \in P'} \max_{g \in A} \left[u(g(\theta)) - \max_{f \in A} u(f(\theta)) \right] \end{aligned}$$

where the second inequality is implied by $co(P) \subseteq co(P')$. But this implies that $h \notin \mathcal{C}'(A, \{h\})$, as desired.

Now prove the second statement. Suppose (u, P) , (u, P') represent \mathcal{C} , \mathcal{C}' , respectively, and suppose that \mathcal{C}' is more diversification averse than \mathcal{C} , and P' satisfies Assumption 2. Let $I = u(\Delta(Z))$; we can without loss assume $[-1, 1] \in I$. We can identify each act f with its corresponding utility act $u \circ f \in I^\Theta$.

Towards a contradiction, suppose that $co(P) \not\subseteq co(P')$; this guarantees the existence of some $p \in co(P) \setminus co(P')$, where $p \in P$. By a separating hyperplane theorem (Dunford and Schwartz, 1957, Theorem V.2.10), there exists a nonzero $\xi \in \mathbb{R}^\Theta$ and $c \in \mathbb{R}$ s.t.

$$\max_{p' \in P'} \xi \cdot p' < c < \xi \cdot p$$

By rescaling, we can without loss assume that $\xi \in I^\Theta$ and $c \in I$. Since P' satisfies Assumption 2, there exists $\epsilon > 0$ such that for acts $(g_\theta)_{\theta \in \Theta}$ satisfying

$$g_\theta(\theta') = \begin{cases} c + \epsilon & \theta' = \theta \\ -1 & \text{otherwise} \end{cases}$$

we have $g_\theta \cdot p' < c$ for all $p' \in P'$ and all θ . There exists $\lambda \in (0, 1)$ such that $(\lambda\xi + (1 - \lambda)c)(\theta) < c + \epsilon$ for all θ ; fix such a λ and let $f = \lambda\xi + (1 - \lambda)c$. Let $A = \{f\} \cup \{g_\theta\}_{\theta \in \Theta}$; by construction, the constant act c is \mathcal{C} -diversified wrt. A .

Since $c < f \cdot p$, $c \notin \mathcal{C}(A \cup \{c\})$. Note, however, since $\max_{p' \in P'} f \cdot p' < c$ and also $\max_{p' \in P'} g_\theta \cdot p' < c$ for all θ , $c = \mathcal{C}'(A \cup \{c\})$. This implies that \mathcal{C}' is not more diversification averse than \mathcal{C} , a contradiction; we have $co(P) \subseteq co(P')$ as desired. \square