

Dynamic Concern for Misspecification*

Giacomo Lanzani - MIT

January 2023 - Click [here](#) for the latest version.

Abstract

We consider an agent who posits a set of probabilistic models for the payoff-relevant outcomes. The agent has a prior over this set but fears the actual model is omitted and hedges against this possibility. The concern for misspecification is endogenous: If a model explains the previous observations well, the concern attenuates. We show that different static preferences under uncertainty (subjective expected utility, maxmin, robust control) arise in the long run, depending on how quickly the agent becomes unsatisfied with unexplained evidence and whether they are misspecified. The misspecification concern's endogeneity naturally induces behavior cycles, and we characterize the limit action frequency. This model is consistent with the empirical evidence on monetary policy cycles and choices in the face of complex tax schedules. Finally, we axiomatize in terms of observable choices this decision criterion and how quickly the agent adjusts their misspecification concern.

,

*I am particularly indebted to Simone Cerreia-Vioglio, Roberto Corrao, Drew Fudenberg, Stephen Morris, and Philipp Strack. I am also grateful to Abhijit Banerjee, Pierpaolo Battigalli, Benjamin Brooks, Yeon-Koo Che, Daniel Clark, Tommaso Denti, Esther Dufo, Federico Echenique, Andrew Ellis, Glenn Ellison, Joel Flynn, Mira Frick, Benjamin Golub, Olivier Gossner, Lars Peter Hansen, Johannes Hörner, Ryota Iijima, Peter Klibanoff, Jacob Leshno, Jonathan Libgober, Fabio Maccheroni, Massimo Marinacci, Moritz Meyer-ter-Vehn, Anh Nguyen, Jawwad Noor, Daisuke Oyama, Parag Pathak, Luciano Pomatto, Giulio Principi, Frank Schilbach, Vasiliki Skreta, Advik Shreekumar, Marciano Siniscalchi, Charles Sprenger, Tomasz Strzalecki, Omer Tamuz, Jean Tirole, Emanuel Vespa, and Alexander Wolitzky for helpful comments and suggestions. The financial support of Guido Cavazzan Scholarship is gratefully acknowledged.

1 Introduction

Bayesian rationality requires that an agent uncertain about the data-generating process postulates multiple probabilistic descriptions of the environment and uses Bayes rule to adjust their relative weights. However, even rational agents may fear that they are misspecified and that none of these descriptions is correct. This concern is remarkably natural in complex and high-dimensional settings, where uncertainty needs to be simplified to obtain well-behaved optimization and learning procedures.

For example, none of the model economies considered by a central bank perfectly describes the underlying data-generating process for output and inflation. Similarly, the consumer response models that a firm uses to set prices and qualities are unlikely to include one that considers all relevant decision factors. Moreover, the diffusion of complex and not explicitly described machine learning algorithms naturally creates new reasons for misspecification. Indeed, consumers increasingly rely on automated recommendations. Although they may have some conjecture on how the alternative’s features translate into a score or a “match quality” with their profile, they certainly do not consider the specific algorithm used by these recommendation systems. Misspecification is even more relevant when dealing with entirely novel issues, such as those faced by a regulatory body that tries to mitigate the effect of climate change using theoretical models that take into account human impacts never experienced in history.

Misspecification has been analyzed from two distinct perspectives. On the one hand, several papers have studied the long-run implications of subjective expected utility (SEU) maximizers learning with misspecified beliefs (see, e.g., Esponda and Pouzo, 2016, Fudenberg, Lanzani, and Strack, 2021, Frick, Iijima, and Ishii, 2022, and the references therein). These works assume that the agents have no concern about being misspecified. Here we show that the absence of such concern is normatively unappealing, as it can induce long-run average payoffs lower than a safe guarantee. It also seems descriptively unrealistic, as the widely documented ambiguity-averse behavior may be seen as a way to hedge against the incorrect specification of the model. On the other hand, the robust control literature in macroeconomics pioneered by Hansen and Sargent (2001) considers agents who fear model misspecification. In particular, the first axioms-based decision criterion that accounts for model misspecification was proposed in Cerreia-Vioglio, Hansen, Maccheroni, and Marinacci (2022).¹

¹It has as a particular case the robust control model of Hansen and Sargent (2001) axiomatized by Strzalecki (2011). Since in Strzalecki (2011) the reference probability is subjective, it can also be

This work reconciles these approaches and shows how popular decision criteria such as maxmin expected utility, robust control preferences, and subjective expected utility arise as the limit behavior of an agent concerned about misspecification and learning about the actual data-generating process (DGP). We consider an agent that repeatedly chooses among actions whose payoffs have an unknown distribution. This choice is taken using an average of robust control assessments, where each assessment takes a different structured model as the benchmark. We introduce endogeneity in the misspecification concern: the better the structured models explain the past, the less concerned the agent is. Specifically, we study agents who repeatedly perform likelihood ratio tests (LRT) of their model and whose concern for misspecification is monotone in the likelihood ratio statistics.

There are two critical determinants for the long-run dynamics: whether the agent is correctly specified and how demanding they are in evaluating their models' performance. First, we consider the case of a correctly specified agent. In that case, the behavior converges to a self-confirming equilibrium, regardless of how demanding the agent is in evaluating their model. A self-confirming equilibrium means that they play an SEU best reply to a belief supported over the data-generating processes that are observationally equivalent to the true one given the chosen action.

Instead, to characterize the limit behavior under misspecification, a taxonomy of how demanding the agent is turns out to be crucial. In particular, a “statistically sophisticated” agent applies a time scaling to the likelihood ratio statistic that keeps it informative about the model's fitness. To support the identification of these statistically sophisticated types with rationality, we show that the achievement of two desirable properties uniquely characterizes them: safety under misspecification (i.e., guaranteeing at least the minmax payoff) and consistency under almost correct specification (i.e., no regret with small misspecification).²

We allow departures from this normative benchmark to obtain descriptive predictions on the effect of an endogenous misspecification concern. We consider agents that are too demanding in evaluating the models' performance (this case includes believers in the Law of Small Numbers, LSN, Tversky and Kahneman, 1971, that treat failures in explaining early realizations as a statistician treats long-run failures). Similarly, we

interpreted as an axiomatization of robust prior analysis, see Hansen and Sargent (2022a).

²Moreover, we observe that SEU maximization, the original robust control of Hansen and Sargent (2001) and the more recent maxmin robust control of Hansen and Sargent (2022b) fail to satisfy these mild rationality requirements jointly.

allow the opposite case in which the agent is too lenient in evaluating their model and attributes too much unexplained evidence to sampling variability.

We then characterize the long-run behavior of these different types of misspecified agents. The actions of the lenient type converge to a Berk-Nash equilibrium, i.e., to an SEU best reply to beliefs supported on the models closest in relative entropy to the actual data-generating process. Instead, overemphasis on the model’s failures in explaining the data by the demanding type induces convergence to a maxmin best reply to the models that are absolutely continuous with respect to the true one.

In contrast, a statistically sophisticated type maintains a non-trivial concern for misspecification. If their behavior converges, it converges to a robust control best reply to the models closest in relative entropy to the actual data-generating process. Moreover, the misspecification concern is endogenously determined by how well the best models fit the evidence generated by the limit action.

Therefore, our learning results provide several novel predictions about the relation between uncertainty attitudes and other individual traits.³ First, the extent of long-run uncertainty aversion positively correlates with the agent being initially misspecified and their belief in the LSN. Second, these correlations are causal: repeated failures to explain the data (misspecification) and demanding evaluation of these failures induce the agent to shift to cautious behavior. Third, even keeping constant the misspecification and understanding of probability rules, the limit uncertainty attitudes are stochastic. Initial realizations leading to a limit action with consequences poorly explained by the agent’s models induce a long-run uncertainty aversion higher than realizations leading to a limit action whose consequences are well explained.

We thus use the equilibrium behavior predicted by an endogenous concern for misspecification to rationalize the labor supply in the face of complex tax schedules documented in Rees-Jones and Taubinsky (2020). In particular, they show that around 40% of the agents have beliefs corresponding to a heuristic that simplifies the tax schedule to a linear one but that 20% fewer agents act accordingly to this heuristic. This is predicted by an endogenous concern for misspecification, as agents with an incorrect model are less prone to base their decisions on the conclusions they reach within the model.

In general, the behavior of a statistically sophisticated type is not guaranteed to converge. Indeed, it is possible that their behavior cycles between phases of different

³The empirical study of the correlation between behavioral biases is an active area of recent development. See, e.g., Dean and Ortleva (2019) and the references therein.

misspecification concerns. Still, we characterize the limit action frequency and concern for misspecification. We apply this result to revisit the cyclical behavior of monetary policies documented in Sargent (1999) and Sargent (2008). Intuitively, the cycles have the following structure. The agent plays an action whose consequences are well explained by one of their structured models (a conservative monetary policy in the application). Playing this action lowers the concern for misspecification and eventually leads to a more misspecification-vulnerable action (a more aggressive monetary policy). Failures to explain the distribution of outcomes observed under this action lead to a return to the more misspecification robust action.

We also obtain two results that provide a testable foundation to the model employed in the learning part of the paper: An axiomatization of the static average robust control criterion and testable axioms for when the agent is of the lenient, statistically sophisticated, or demanding type. Two primary axioms pin down the static decision criterion. The first is a weaker form of the Sure-Thing Principle imposed only on bets on the data-generating process (e.g., bets on the urn composition) and bets conditional on the data-generating process (e.g., bets on the ball color conditional on having been told the urn composition). The second requires that conditional on being told the best-fitting model, the agent is equally concerned about misspecification regardless of which one it is.

For the dynamic representation, a dynamic consistency axiom on the acts that bet on the data-generating process is shown to guarantee Bayesian updating over models. More interestingly, the preference adjustment of a statistically sophisticated type is pinned down by a novel Asymptotic Frequentism axiom, requiring arbitrarily similar preferences conditional to sufficiently long histories with the same outcome frequency.

The rest of the paper is structured as follows. Section 2 introduces the average robust control decision criterion and how preferences are adjusted. Section 3 studies what attitudes toward model failures induce good payoff performance and provides a learning foundation for the different uncertainty attitudes. Section 4 characterizes the limit frequency of time spent using the different actions when the behavior does not converge and applies the result to a central banking problem. Section 5 provides the axiomatization to the static decision criterion and how preferences are updated. Section 6 discusses the related literature and possible extensions. Section 7 concludes. All proofs are collected in the Appendix.

2 Decision Criterion

2.1 Static Decision Criterion

We describe the criterion used in the repeated decision problem and defer its axiomatization to Section 5. We consider an agent who evaluates a finite number of actions $a \in A$ and let Y be a compact subset of a Euclidean space representing the set of possible outcomes. The agent has a continuous utility index $u : A \times Y \rightarrow \mathbb{R}$ over the action-outcome pairs that captures their preference when the subjective uncertainty is resolved. However, the realized outcome is stochastic and endogenous as each action $a \in A$ induces an objective probability measure $p_a^* \in \Delta(Y)$ over outcomes.⁴

Subjective Beliefs The agent correctly believes that the map from actions to probability distributions over outcomes is fixed and depends only on their current action. Still, they do not know $p^* = (p_a^*)_{a \in A}$ and deal with this uncertainty in a quasi-Bayesian way. The agent postulates a set of *structured models*, i.e., action-dependent probability measures over outcomes $q^\theta = (q_a^\theta)_{a \in A} \in \Delta(Y)^A$, indexed by a compact and finite-dimensional set of parameters Θ . They have a prior belief μ with support Θ that describes the relative likelihood assigned to these models. For example, the agent may be a central bank that considers a Keynesian Samuelson-Solow model where the monetary policy affects the unemployment rate or a new classical Lucas-Sargent model with no systematic effect of inflation on unemployment.⁵

We must impose a few regularity conditions.

Assumption 1. For every $a \in A$ and p_a^* -almost every $y \in Y$: (i) For all $\theta \in \Theta$, $p_a^* \sim q_a^\theta$ and the density of q_a^θ with respect to p_a^* , denoted as \tilde{q}_a^θ , is continuous and p_a^* -a.s. bounded away from 0, uniformly over θ ,⁶ (ii) The map $\theta \mapsto \tilde{q}_a^\theta(y)$ is continuous.

Condition (i) allows us to compute the relevant expectations while allowing for both discrete and continuous outcome spaces. Continuity of the map from parameters to

⁴For every subset C of a Euclidean space, we denote as $\Delta(C)$ the Borel probability measures on C , endowed with the topology of weak convergence of measures.

⁵This formulation follows the recent literature on misspecified learning in assuming that both the true data-generating process and the subjective models considered by the agent are i.i.d. conditionally on the agent's behavior. This makes the true extent of misspecification time-invariant and "learnable". It is important to notice that this time invariance is often relaxed in the literature on dynamic decisions with robust control preferences that follows Hansen and Sargent (2001).

⁶For every $p, q \in \Delta(Y)$, $p \gg q$ means that q is absolutely continuous with respect to p , and $p \sim q$ means that they are mutually absolutely continuous.

outcome distributions is a standard requirement for parametric models. Condition (iii) guarantees that no subjective model of the agent is ruled out in finite time.⁷

A Bayesian agent with complete trust in their set of models evaluates an action $a \in A$ according to its subjective expected utility (see, e.g., Cerreia-Vioglio, Maccheroni, Marinacci, and Montrucchio, 2013b):

$$\int_{\Theta} \mathbb{E}_{q_a^\theta} [u(a, y)] d\mu(\theta).$$

That is, they compute a two-stage expectation of the utility function: they evaluate the utility of the action given the candidate model θ , $\mathbb{E}_{q_a^\theta} [u(a, y)]$, and then they average over the models with weights given by their subjective belief μ .

However, we are interested in agents concerned with the possibility that none of these models is the exact description of the data-generating process but only a valid approximation, i.e., that are concerned that there is no $\theta \in \Theta$ with $q^\theta = p^*$. Therefore, in the spirit of the robustness criterion advocated by Hansen and Sargent (2001), they penalize actions that perform poorly under alternative distributions that are close in relative entropy $R(\cdot||\cdot)$ to some of the structured models.⁸

With this, an agent evaluates each action $a \in A$ accordingly to the *average robust control* criterion:

$$\int_{\Theta} \min_{p_a \in \Delta(Y)} \mathbb{E}_{p_a} [u(a, y)] + \frac{1}{\lambda} R(p_a || q_a^\theta) d\mu(\theta) \quad (1)$$

where $\lambda > 0$ is a parameter that trade-offs between decision robustness and performance under the structured models.⁹

The standard robust control model introduced by Hansen and Sargent (2001) is the case in which μ is a Dirac measure (that in macroeconomics applications is often assumed to satisfy rational expectations, i.e., to be degenerate on the actual data-generating process). As described in Hansen, Sargent, Turmuhambetova, and Williams (2006), this case corresponds to when “[...] a maximizing player (‘the decision maker’)

⁷Part (iii) also plays a technical role in guaranteeing the existence of the equilibrium concepts we consider. It is known it can be relaxed, see Anderson, Duanmu, Ghosh, and Khan (2022), but this relaxation comes at the cost of requiring nonstandard analysis techniques (where nonstandard means using infinitesimal numbers), something beyond this paper’s scope.

⁸Recall that for every $p, p' \in \Delta(Y)$, $R(p||p') = \int_Y \log\left(\frac{dp}{dp'}\right) dp$ if $p' \gg p$ and $R(p||p') = \infty$ otherwise. Appendix A.3 explains under what other distances between probability distributions our results continue to hold.

⁹Lemma 1 justifies the use of a min rather than an inf in equation (1) and throughout the paper.

chooses a best response to a malevolent player (‘nature’) who can alter the stochastic process within prescribed limits. The minimizing player’s malevolence is the maximizing player’s tool for analyzing the fragility of alternative decision rules.” Equation (1) extends this interpretation to a situation in which the agent is still uncertain about the best-approximating model (i.e., μ is nondegenerate), allowing the malevolent nature to alter each of the candidate structured models.

The representation adopts the distinction between two levels of uncertainty. At the first level, given a probabilistic model q^θ , the uncertainty about the exact specification of the model is captured by minimizing the expected utility for probabilities that are not too far away from q^θ . At a higher level, the agent is also uncertain about the identity of the best structured model and posits a prior probability μ over them. While the higher level of uncertainty is already present under subjective expected utility, the lower level captures the agent’s concern for misspecification.

2.2 Preference Evolution

The average robust control criterion of equation (1) describes how the agent chooses for a *given* belief and level of misspecification concern. However, preferences respond to the received information.

On the one hand, we stick to the classical dynamic treatment of tastes over certain alternatives and beliefs about the possible data-generating processes. We let the utility index u be constant over time, and the belief be updated through standard Bayesian updating.

Formally, time is discrete, and a history is a finite vector of past actions and outcomes. In particular, the set of histories of finite length $t \in \mathbb{N}$ is $\mathcal{H}_t = (A \times Y)^t$ and the set of all finite histories is $\mathcal{H} = \bigcup_{t=0}^{\infty} \mathcal{H}_t$. With this, for every measurable subset C of Θ , we denote by

$$\mu(C \mid (a^t, y^t)) = \frac{\int_{\theta \in C} \prod_{\tau=1}^t \tilde{q}_{a_\tau}^\theta(y_\tau) d\mu_0(\theta)}{\int_{\theta \in \Theta} \prod_{\tau=1}^t \tilde{q}_{a_\tau}^\theta(y_\tau) d\mu_0(\theta)} \quad (\text{Bayes Rule})$$

the subjective belief the agent obtains using Bayes rule after history $(a^t, y^t) \in \mathcal{H}_t$.¹⁰

On the other hand, we introduce an endogenous and time-evolving concern for misspecification, i.e., λ depends on the realized history. In particular, we want the concern

¹⁰By Assumption 1 (ii), the posterior is well-defined after every positive probability history. We allow for arbitrary belief revisions after events with zero ex-ante subjective probability.

for misspecification to be a function of how well the structured models explain the current history. In statistics, the most standard measure of fit of a set of distributions Θ against a set of unstructured alternatives $N(\Theta) \subseteq \Delta(Y)^A$ is the likelihood ratio test statistic:¹¹

$$LRT((a^t, y^t), \Theta) = -2 \log \left(\frac{\max_{\theta \in \Theta} \prod_{\tau=1}^t \tilde{q}_{a_\tau}^\theta(y_\tau)}{\max_{p \in N(\Theta)} \prod_{\tau=1}^t \tilde{p}_{a_\tau}(y_\tau)} \right) \quad \forall t \in \mathbb{N}, \forall (a^t, y^t) \in \mathcal{H}_t.$$

Here we want to take a conservative approach and not impose structure over the set of alternative unstructured distributions $N(\Theta)$ used to evaluate the model's fit. If Y is finite (or, under some regularity conditions, countable) and all outcomes have positive probability, there is a natural way to do so, i.e., to consider as the set of alternatives unstructured distributions the entire (action-indexed) simplex $\Delta(Y)^A$. However, considering a completely unrestricted set of distributions with a continuum of outcomes leads to an utterly uninformative test of the model, as the (discrete) empirical distribution is an infinitely better fit to itself than any continuous distribution, i.e., the test always returns $+\infty$. To maintain informativeness, $N(\Theta)$ must then include only distributions that are mutually absolutely continuous with respect to the ones in Θ . In particular, all our results are invariant to the $N(\Theta)$ choice as long as the following assumption is satisfied.

Assumption 2. (i) $N(\Theta) \supseteq \Theta$ is closed and $p^* \in N(\Theta)$. (ii) The family of densities $(\tilde{p}_a)_{p \in N(\Theta)}$ is equicontinuous.

We require that the unstructured set is a relaxation of the parametric structure sufficiently large to include the actual distribution and a continuity condition that rules out a Θ that only contains continuous distributions and an $N(\Theta)$ that includes discrete distributions.

With this, we let the updated concern for misspecification to be

$$\lambda^\beta(h_t) = \frac{LRT(h_t, \Theta)}{2\beta_t} \quad \forall t \in \mathbb{N}, \forall h_t \in \mathcal{H}_t \quad (2)$$

where $\beta = (\beta_t)_{t \in \mathbb{N}} \in \mathbb{R}_{++}^\mathbb{N}$ is an unbounded increasing sequence. Of course, we do not

¹¹The Neyman-Pearson Lemma establishes the performance of the test under correct specification. At the same time, Foutz and Srivastava (1977) and Vuong (1989) contain the classical results about the informativeness of the LRT under misspecification. Schwartzstein and Sunderam (2021) is a recent paper that models agents in a persuasion problem who perform model selection using this statistic. Lemma 1 justifies the use of max in the definition of the LRT under Assumption 2.

commit to a literal interpretation of agents performing an LRT of their model: we only require that they behave in a more misspecification-concerned way after histories for which their model fit is worse.¹²

Beyond the likelihood ratio test, a key role is played by the sequence of time scalings β . Indeed, if the agent is misspecified, the expected one-period increase in the LRT is strictly positive, regardless of the distance between the actual DGP and the set of models in Θ . Therefore, if we want to capture agents that adjust their misspecification concern $\lambda^\beta(h_t)$ in a way that remains responsive to the model fit, and that does not mechanically and systematically increase over time, some scaling must be applied.

In particular, the statistics literature shows that a time scaling of the form $\beta_t = ct$ for some $c \in \mathbb{R}$ is the only one that gives a statistic whose asymptotic distribution is increasing in the distance between Θ and the true DGP; see, e.g., Foutz and Srivastava (1977) and Vuong (1989). For this reason, the likelihood ratio with linear time scaling is used to measure the extent of model misspecification.¹³ Motivated by these results, we often informally refer to an agent who uses such a time scaling as a “statistically sophisticated type”. Of course, the objective of a statistician can be very different from that of an agent involved in a decision problem under uncertainty. Proposition 1 confirms that this time scaling is also a rationality benchmark in repeated decision problems, as it uniquely identifies the behavior that induces no regret when the agent is correctly specified and is always maxmin safe.

Moreover, when $\beta_t = ct$, the combination of Bayesian updating over parameters and dynamically adjusted concern for misspecification is consistent with Savage’s distinction between small and large worlds (see pages 82-91 in Savage, 1954). Indeed, Savage advocates reducing the large-world uncertainty to small worlds (for us, the finite-dimensional Θ) where Bayesian updating has appealing properties, but being aware that this description is incomplete and that the agent should evaluate the fit of that simplification (for us, using a test that can measure the failures of this description).

At the same time, we capture some less normatively appealing but descriptively relevant phenomena by allowing nonlinear time scaling. On the one hand, a time

¹²We discuss in Section 6.3 the most general form of dependence on model fit under which our results continue to hold.

¹³This use of the LRT complements its classical role in deciding whether to reject or accept a model. In particular, Wilks’ Theorem (see, e.g., Theorem 10.3.3 in Casella and Berger, 2021) shows that *under correct specification*, the likelihood ratio statistic converges to a χ^2 distribution. However, it says nothing about the distribution of the LRT if the model is misspecified. See Hausman (1978) and the subsequent literature for a complementary approach to the measurement of model misspecification when the statistician can compute a consistent quasi-maximum-likelihood estimator.

scaling such that $\beta_t = o(t)$, e.g., $\beta_t = \sqrt{t}$, overly penalizes minor imperfections of the model, expecting that the frequency quickly converges to its theoretical value, as in the fallacy called the Law of Small Numbers. On the other hand, an agent for which $t = o(\beta_t)$ applies an excessively lenient time scaling to the likelihood ratio statistic and attributes too much of the unexplained evidence to sampling variability.

3 Long-run Payoffs and Actions

In this section, we study the long-run consequences of using the decision criterion and the preference updating rule proposed above. Our primary interest is in what attitudes towards model failures induce good payoff performance across environments and what are the limit actions and preferences under uncertainty attitudes that arise given a specific attitude.

Let $BR^\lambda(\nu)$ denote the set of average robust control best replies to belief ν when the concern for misspecification is λ , i.e.,¹⁴

$$BR^\lambda(\nu) = \operatorname{argmax}_{a \in A} \int_{\Theta} \min_{p_a \in \Delta(Y)} \mathbb{E}_{p_a} [u(a, y)] + \frac{1}{\lambda} R(p_a || q_a^\theta) d\nu(\theta) \quad \forall \lambda \geq 0, \forall \nu \in \Delta(\Theta).$$

Also let

$$BR^{Seu}(\nu) = \operatorname{argmax}_{a \in A} \int_{\Theta} \mathbb{E}_{q_a^\theta} [u(a, y)] d\nu(\theta) \quad \forall \nu \in \Delta(\Theta)$$

denote the actions that maximize the (classical) subjective expected utility of an agent with belief ν and

$$BR^{Meu}(C) = \operatorname{argmax}_{a \in A} \inf_{p \in C} \mathbb{E}_{p_a} [u(a, y)]$$

denote the actions preferred by a maxmin agent a la Gilboa and Schmeidler (1989) with a set of models $C \subseteq \Delta(Y)^A$.

A (pure) *policy* is a measurable $\Pi : \mathcal{H} \rightarrow A$ that specifies an action for every history. The objective action-contingent probability distribution and a policy Π induce a probability measure \mathbb{P}_Π on $(A \times Y)^\mathbb{N}$.¹⁵ Our interest is in policies derived from maximizing the value in equation (1) for some rule β determining how the concern for misspecification is adjusted.

Definition 1. Policy Π is β -optimal if for all $h_t \in \mathcal{H}$, $\Pi(h_t) \in BR^{\lambda^\beta(h_t)}(\mu(\cdot|h_t))$.

¹⁴Throughout the paper, we use the convention $0 \cdot \infty = 0$.

¹⁵We spell out the \mathbb{P}_Π derivation in Appendix A.1.1.

3.1 Safety and Consistency

We have mentioned that using a linear time scaling β has the good statistical property of keeping λ^β asymptotically informative about the fit of the model. Now, we provide a normative justification for considering it the relevant benchmark of rationality, showing that it satisfies the desirable properties of safety and consistency (cf. Fudenberg and Levine, 1995) across all possible exogenous decision problems the agent can face.

Definition 2. Let $\varepsilon > 0$. β is ε -safe for the decision problem (u, A, Y) if for every β -optimal policy Π and DGP $p^* \in \Delta(Y)^A$

$$\liminf_{t \rightarrow \infty} \frac{\sum_{i=1}^t u(a_i, y_i)}{t} \geq \max_{a \in A} \min_{y \in Y} u(a, y) - \varepsilon \quad \mathbb{P}_\Pi\text{-a.s.} \quad (3)$$

This is a very mild condition that only requires the agent to obtain an average payoff at least ε close to what they can guarantee against *every* possible outcome. However, when paired with misspecification, ε -safety has a significant bite: a Bayesian SEU agent fails it in many decision problems. Indeed, such failures have been the basis of many critiques of learning under misspecification with Bayesian SEU agents (see, e.g., Massari and Newton, 2020).

We say that a decision problem (u, A, Y) is *exogenous* if $q_a^\theta = q_{a'}^\theta$ and $p_a^* = p_{a'}^*$ for all $\theta \in \Theta$ and $a, a' \in A$. Exogenous decision problems are situations in which the agent correctly believes that the action does not affect the outcome distribution.

Definition 3. Let $\varepsilon > 0$. β is ε -consistent under almost correct specification for the exogenous decision problem (u, A, Y) if there exists $\delta > 0$ such that for every β -optimal policy Π and DGP $p^* \in \Delta(Y)$

$$\min_{\theta \in \Theta} R(p^* || q^\theta) < \delta \implies \liminf_{t \rightarrow \infty} \frac{\sum_{i=1}^t u(a_i, y_i)}{t} \geq \max_{a \in A} \mathbb{E}_{p^*} [u(a, y)] - \varepsilon \quad \mathbb{P}_\Pi\text{-a.s.}$$

ε -consistency under almost correct specification requires that sufficiently low levels of misspecification (i.e., the existence of a model q^θ with distance less than δ from the true data generating process) cannot induce considerable ex-post regret (i.e., a limit average payoff more than ε lower than the expected payoff of the objectively optimal action). Intuitively, we want that if the misspecification is minor, in the long run, the agent approximately identifies the actual model and starts best replying to it.

Proposition 1. 1. For every exogenous decision problem and $\varepsilon > 0$ there exists $c > 0$ such that $\beta_t = ct$ is both ε -safe and ε -consistent under correct specification.

2. *There exists an exogenous decision problem and $\varepsilon > 0$ for which there is no ε -safe and ε -consistent under almost correct specification β with either $\beta_t = o(t)$ or $o(\beta_t) = t$.*

The safety and consistency conditions we require are weak but are enough to single out the statistically sophisticated type. Moreover, even if, for simplicity of exposition, the statement of the proposition maintains the ratio structure of equation (2), in the Appendix, we prove the stronger claim that there are decision problems where any way to adjust the concern for misspecification that is globally more demanding or lenient than the average LRT violates either ε -safety or ε -consistency under almost correct specification. In contrast, standard SEU maximization is not safe, while always using a maxmin best reply to $(q_a^\theta)_{\theta \in \Theta}$ induces a behavior that is safe but not consistent under correct specification.¹⁶

3.2 Long-run Behavior

We are interested in the actions that can arise as the long-run behavior of agents with an evolving concern for misspecification. The main results of this section show that we can describe this limit behavior through fixed point conditions involving the agent's action, belief, and concern for misspecification. To this end, let $\Theta(a) = \operatorname{argmin}_{\theta \in \Theta} R(p_a^* || q_a^\theta)$ be the structured models that best fit the actual data-generating process when action a is played.

Definition 4. Action a^* is a:

1. *Self-confirming equilibrium* (SCE) if there exists $\nu \in \Delta(\Theta)$ with

$$\operatorname{supp} \nu \subseteq \{\theta \in \Theta : q_{a^*}^\theta = p_{a^*}^*\} \text{ and } a^* \in BR^{Seu}(\nu).$$

2. *Berk-Nash equilibrium* (B-NE) if there exists $\nu \in \Delta(\Theta)$ with

$$\operatorname{supp} \nu \subseteq \Theta(a^*) \text{ and } a^* \in BR^{Seu}(\nu).$$

¹⁶This observation about the inconsistency of a misspecified single agent complements the result of Fudenberg and Kreps (1993) and Fudenberg and Levine (1995) about the inconsistency of a correctly specified SEU player in games.

3. *Maxmin equilibrium* if

$$a^* \in BR^{Meu} \left(\left\{ p \in \Delta(Y)^A : \exists \theta \in \Theta, \forall a \in A, q_a^\theta \gg p_a \right\} \right).$$

4. *c-robust equilibrium* if there exists $\nu \in \Delta(\Theta)$ with

$$\text{supp } \nu \subseteq \Theta(a^*), \quad a^* \in BR^\lambda(\nu), \quad \text{and } \lambda = \min_{\theta \in \Theta} R(p_{a^*}^* || q_{a^*}^\theta) / c.$$

Self-confirming equilibrium (Battigalli, 1987 and Fudenberg and Levine, 1993) describes a stable situation where the agent's action is a best reply to a belief that is on-path confirmed, in the sense of being concentrated over models that perfectly match the distribution over outcomes induced by the equilibrium action.

Berk-Nash equilibrium (Esponda and Pouzo, 2016) relaxes the confirmed beliefs condition of SCE by only requiring that the supporting beliefs are concentrated on the models that provide the best fit to the outcome distribution induced by the equilibrium action. Importantly, this fit is not required to be perfect.

In a maxmin equilibrium, the agent evaluates each action under the worst-case scenario that is minimally consistent with their structured descriptions of the environment (i.e., those scenarios that do not assign positive probability to events impossible for the structured models).

c-robust equilibrium is similar to Berk-Nash in requiring best reply to the best-fitting models. However, the best reply is the average robust control, with misspecification concern that decreases in how well the models fit the true DGP at the equilibrium.

We are interested in what actions have a positive probability of becoming the long-run behavior of the agent. The following definition captures this requirement.

Definition 5. Action a is a β -limit action if there is a β -optimal policy Π such that $\mathbb{P}_\Pi[\sup\{t: a_t \neq a\} < \infty] > 0$.

Our first limit result is a consistency check: Concern for misspecification is irrelevant in environments with a finite number of outcomes if the agent is correctly specified about the consequence of the limit action.

Proposition 2. If Y is finite and a^* is a β -limit action with $p_{a^*}^* \in \text{int} \{q_{a^*}^\theta\}_{\theta \in \Theta}$ then a^* is a self-confirming equilibrium.

Instead, how quickly the agent becomes unsatisfied with their model plays a key role when misspecified.

Theorem 1. *Let a^* be a β -limit action with $p_{a^*}^* \notin \{q_{a^*}^\theta\}_{\theta \in \Theta}$. We have:*

1. *If $t = o(\beta_t)$, then a^* is a Berk-Nash equilibrium.*
2. *If $\beta_t = o(t)$, then a^* is a maxmin equilibrium.*
3. *If $\beta_t = ct$, then a^* is a c -robust equilibrium.*

The theorem characterizes the possible limit actions of all types of agents. At one extreme, the concept of Berk-Nash equilibrium, introduced for subjective expected utility maximizers, is still sufficient to describe the long-run behavior of lenient types. At the other extreme, the repeated failures in explaining the observed data lead demanding agents to a highly pessimistic behavior and consider the worst-case scenario among all the DGPs that are minimally consistent with the structured models.

Finally, if the behavior of the statistically sophisticated type converges, the limit action a^* is a best reply to beliefs that are supported on the relative entropy minimizers. Here the misspecification concern is determined by the relative entropy between the actual DGP and the best-fitting model.

3.3 Equilibrium Illustrations

In this section, we revisit two of the main biases that have been justified as a consequence of misspecified learning (see Esponda and Pouzo, 2016). Within each example, adding an endogenous concern for misspecification predicts a change in a clear direction. However, one bias is reduced while the other is enhanced. Both changes are broadly consistent with the documented evidence. The first example shows how the endogenous misspecification concern moderates the Berk-Nash equilibrium's prediction that a more complicated tax schedule induces a higher labor supply.

Example 1 (Bias Reduction under Misperceived Taxation, Sobel, 1984 and Esponda and Pouzo, 2016). *An agent chooses effort $a \in A$ at cost $c(a)$ and obtains income $z = a + \omega_a$, where ω_a is a stochastic term with $\mathbb{E}_{p_a^*}[\omega_a] = 0$ for all $a \in A$. The agent pays taxes $t = \tau(z) + \varepsilon_1$, where $\tau : \mathbb{R} \rightarrow \mathbb{R}$ is a convex tax schedule. Here $y = (z, t)$, and the payoff is $u(a, y) = z - t - c(a)$. The agent believes in a random coefficient model, $t = (\theta + \varepsilon_2)z$, in which the marginal and average tax rates are both equal to $(\theta + \varepsilon_2)$ and $\Theta \subseteq \mathbb{R}$. The stochastic terms $\varepsilon_1, \varepsilon_2 \sim N(0, 1)$ measure respectively actual and conjectured uncertain aspects of the tax schedule, and the $(\omega_a)_{a \in A}$, ε_1 , and ε_2 are*

independent.¹⁷ See Liebman and Zeckhauser (2004) and Rees-Jones and Taubinsky (2020) for the empirical evidence supporting this “schmeduling” bias.

Simple computations show that $\Theta(a) = \left\{ \mathbb{E}_{p_a^*} \left[\frac{\tau(a+\omega_a)}{a+\omega_a} \right] \right\}$, i.e., the best fitting marginal taxation is equal to the (lower) average taxation.¹⁸ Therefore, as pointed out by Esponda and Pouzo (2016), in any Berk-Nash equilibrium, the agent ends up exerting higher effort than the optimal. Moreover, the more complex (i.e., convex) the tax code is, the more significant the gap between the average and marginal rate and the higher the excess effort of the agent.

In every c -robust equilibrium, this bias is reduced. To see this observe that since the agent is not perfectly able to explain the equilibrium data, i.e., $\min_{\theta \in \Theta} R(p_a^* || q_a^\theta) > 0$, they maintain a positive level of concern for misspecification. However, higher efforts are perceived as more exposed to the uncertainty in the marginal rate (as the stochastic term $\theta + \varepsilon$ gets multiplied by an, on average, higher z).

Therefore, c -robust equilibrium provides a natural force that reduces the counterintuitive prediction that complicated nonlinear taxation codes induce more effort: failures to rationalize the received tax bill reduce effort. Moreover, the more complicated the tax code is, i.e., the more nonlinear τ is, the larger the correction size. This set of predictions is consistent with Rees-Jones and Taubinsky (2020), where it is shown that around 40% of the agents have beliefs (elicited in an incentive-compatible way) corresponding with the schmeduling heuristic but that there are 20% fewer agents who act accordingly to the schmeduling heuristic.¹⁹

The second example shows that an endogenous concern for misspecification can enhance some biases. In particular, this is the case for Correlation Neglect, a bias that is indeed widely documented (see Enke and Zimmermann, 2019 and the references therein).

Example 2 (Bias Increase under Correlation Neglect, Esponda, 2008). A buyer with valuation $v \in V$ and a seller submit a (bid) price $a \in A$, and an ask price $s \in S \subseteq \mathbb{R}_+$,

¹⁷Formally, ε normally distributed implies that Y is not compact, in contrast with the primary analysis of the paper. Still, the conclusions below are unaffected by considering ε with a symmetrically truncated normal distribution that allows remaining in our main framework.

¹⁸See Appendix A.4 for the computations supporting the claims of the examples.

¹⁹In this discussion we followed Rees-Jones and Taubinsky (2020) preferred interpretation in terms of an heterogeneous population. They observe that their data are also compatible with all the agents having beliefs induced by the schmeduling heuristic but under-responding to this biased estimation of the marginal tax rate. This explanation is consistent with a c -robust equilibrium and inconsistent with a Berk-Nash equilibrium, too.

respectively. They play a double auction with price at the buyer's bid, so the seller sets their ask s equal to their value, and a sale occurs if the buyer's bid a is at least s . The payoff for the buyer is

$$u(a, v, s) = \begin{cases} v - a & a \geq s \\ 0 & \text{otherwise.} \end{cases}$$

The buyer mistakenly believes that the ask price and valuation are independent: $q_a^\theta = \theta$ and $\Theta = \Delta(V) \times \Delta(S)$. Easy computations show that

$$\Theta(a^*) = \left\{ \theta \in \Theta : \forall a \in A, \forall s \in S, q_a^\theta(s) = p_a^*(s), q_a^\theta(v) = p^*(v) \right\}.$$

Therefore, in the Berk-Nash equilibrium, the agent makes a bid a^* lower than the optimal one, not realizing that higher successful bids are, on average, associated with higher quality goods. In this case, the bias is reinforced in a c -robust equilibrium: a complete unraveling of the market where the buyer bids 0 is easier to achieve. The correlation between valuations and prices results in a positive $\min_{\theta \in \Theta} R(p_{a^*}^* || q_{a^*}^\theta) > 0$ and makes the agent less confident in their model. Since offering 0 gives a certain payoff, it is less sensitive to the model misspecification, and, therefore, this positive concern makes market participation less desirable.

4 Cycles

Part 3 of Theorem 1 provides a necessary condition for the limit actions of the statistically sophisticated type. However, as momentarily illustrated by the monetary policy application of Section 4.1, there is no guarantee that such an action exists. In these cases, we know by Theorem 1 that the agent behavior cannot stabilize. We now propose a generalization of c -robust equilibrium, show that it always exists, and prove that it characterizes a weaker form of behavior convergence. Formally, for every $\alpha \in \Delta(A)$, let

$$\Theta(\alpha) = \operatorname{argmin}_{\theta \in \Theta} \sum_{a \in A} \alpha(a) R(p_a^* || q_a^\theta)$$

be the set of parameters with the lowest average relative entropy from the actual data-generating process, where the average is computed using α .

Definition 6. A mixed action $\alpha^* \in \Delta(A)$ is a *mixed c -robust equilibrium* if there

exists $\nu \in \Delta(\Theta)$ with

$$\text{supp} \nu \subseteq \Theta(\alpha^*), \alpha^* \in \Delta(BR^\lambda(\nu)), \text{ and } \lambda = \min_{\theta \in \Theta} \sum_{a \in A} \alpha^*(a) R(p_a^* || q_a^\theta) / c.$$

A mixed robust equilibrium allows multiple actions to be played but requires that the beliefs and the concern for misspecification are determined by the probability assigned to each action. Intuitively, suppose actions for which the models in Θ do not satisfactorily explain the consequences are played more often. In that case, the mixed action α^* must best reply to a more significant misspecification concern.

Proposition 3. *For every $c > 0$ there exists a mixed c -robust equilibrium.*

Existence is established by proving that the conditions characterizing a mixed c -robust (single-agent) equilibrium are equivalent to the ones of a Nash equilibrium in a game among the agent and two adversarial Nature players. The result is then obtained by showing that this game satisfies the conditions that guarantee existence in Reny (1999).

Theorem 1 assumes convergence and characterizes the possible limit actions. However, there are natural environments where the action process almost surely does not converge. In that case, it is important to study a weaker form of behavior stabilization, i.e., the convergence of the empirical distribution over actions, that allows for persistent changes in actions and misspecification concerns. Let $\alpha_t(h_t) \in \Delta(A)$ be the empirical action frequency in history h_t , defined as

$$\alpha_t(h_t)(a) = \frac{\sum_{\tau=1}^t \mathbb{I}_{\{a\}}(a_\tau)}{t} \quad \forall a \in A, \forall h_t \in \mathcal{H}_t.$$

Definition 7. A mixed action $\alpha \in \Delta(A)$ is a β -limit frequency if there is a β -optimal policy Π such that $\mathbb{P}_\Pi[\lim_{t \rightarrow \infty} \alpha_t(h_t) = \alpha] > 0$.

The following result shows that mixed robust equilibria are the relevant equilibrium concept to capture the long-run stabilization of the average time spent playing each action.

Theorem 2. *If $\beta_t = ct$ for all $t \in \mathbb{N}$ and α^* is a β -limit frequency, then α^* is a mixed c -robust equilibrium.*

To interpret Theorem 2, consider the case where α^* is supported over two actions a, a' such that Θ explains very well the consequences of a , —i.e., $\min_{\theta \in \Theta} R(p_a^* || q_a^\theta)$ is

low— but it explains poorly the consequences of a' —i.e., $\min_{\theta \in \Theta} R(p_{a'}^* || q_{a'}^\theta)$ is high. Suppose also that a is a best reply to a high misspecification concern, while a' is a best reply to a low misspecification concern. Then, the agent oscillates between periods with great concern for misspecification, when they play a , and phases in which the excellent data fit leads them to experiment with action a' .

Whenever cycles are involved, a natural concern is whether the agent can predict them and whether they have the incentive to break them.²⁰ This is not the case in this model for two orders of reasons. First, the oscillations in behavior are stochastic, and the agent cannot predict and anticipate the changes perfectly. Second and more important, although the agent behavior does not converge, whenever $\Theta(\alpha)$ is a singleton, the agent's preferences converge.²¹ They are approaching indifference between all the actions with positive frequency. This asymptotic indifference dramatically reduces the incentives to try to detect the probabilistic cycle and break them.

4.1 Application: Monetary Policy Cycles

Here we consider a monetary policy model taken from Sargent (1999), Cogley and Sargent (2005), and Sargent (2008) and in particular its adaptation in Battigalli, Cerreia-Vioglio, Maccheroni, Marinacci, and Sargent (2022). A central bank is trying to control a two-dimensional consequence $Y \subseteq \mathbb{R}^2$, where the y_U component is unemployment and the y_π component is inflation. The policy is aggressive $a = 1$ or conservative $a = 0$.²² The central bank set of models is parametrized by the vector θ , with the following specification:

$$\begin{aligned} y_U &= \theta_0 + \theta_{1\pi} y_\pi + \theta_{1a} a + \theta_2 \varepsilon_U \\ y_\pi &= a + \theta_3 \varepsilon_\pi \end{aligned}$$

where ε_U and ε_π are independent, zero-mean random shocks normalized to have the same support $[-1, 1]$. Here $\theta_0 > 0$ is the natural unemployment level, $\theta_{1\pi} < 0$ is the impact of the actual inflation on unemployment, and $\theta_{1a} > 0$ is the impact of the planned inflation on unemployment, a reduced form of the fact that the market participants (partially) incorporate the central bank actions in their inflation expectations. In par-

²⁰For example, this is what happens under fictitious play.

²¹A singleton $\Theta(\alpha)$ is a mild requirement satisfied in many cases. See Fudenberg, Lanzani, and Strack (2021) for a discussion.

²²Two actions are assumed for simplicity, but a finite A is needed to apply our results.

ticular, if $\theta_{1\pi} + \theta_{1a} = 0$, this is a Lucas-Sargent model with no (structural) exploitable employment-inflation trade-off. If $\theta_{1\pi} + \theta_{1a}$ is negative, this is a Samuelson-Solow model with a structural exploitable employment-inflation trade-off.

The agent's model is misspecified in that it misses the fact that an aggressive monetary policy, beyond raising its baseline level, also increases the inflation variability:

$$\begin{aligned} y_U &= \theta_0^* + \theta_{1\pi}^* y_\pi + \theta_{1a}^* a + \theta_2^* \varepsilon_U \\ y_\pi &= a + \theta_3^* f_a(\varepsilon_\pi) \end{aligned}$$

where f_0 is the identity function, while f_1 is a continuous, strictly increasing, and odd function with $f_1(1) = 1$ that is strictly concave on \mathbb{R}_{++} , i.e., that amplifies the inflation-specific shocks.²³ This form of misspecification is motivated by the findings in Primiceri (2005) and Sims and Zha (2006) and recent inflation consequences of an aggressive monetary policy.

The central bank is endowed with standard quadratic preferences:

$$u(a, (y_U, y_\pi)) = -y_U^2 - y_\pi^2.$$

Assumption 3. i) Some trade-off is present: $\theta_{1\pi}^* + \theta_{1a}^* < -1$. ii) Inflation is more volatile than unemployment under the aggressive monetary policy: $\text{essinf}_{p_1^*} u(1, y) < \text{essinf}_{p_0^*} u(0, y)$. iii) Θ is a product set that includes θ^* and for all $\theta \in \Theta$, $(\theta_{1a}, \theta_2, \theta_3) = (\theta_{1a}^*, \theta_2^*, \theta_3^*)$.

Observe that the exploitable trade-off required by (i) may be so small that the reduced inflation variability under a conservative policy makes the latter optimal. Condition (ii) requires that the additional inflation volatility induced by the aggressive policy is enough to have the worst tail payoffs. Condition (iii) allows us to focus on the cycles induced by the oscillation in the concern for misspecification. Without that, one would get the same insights with other oscillations of beliefs that push even more

²³An alternative, more parametric specification would have

$$\begin{aligned} y_U &= \theta_0^* + \theta_{1\pi}^* y_\pi + \theta_{1a}^* a + \theta_2^* \varepsilon_U \\ y_\pi &= (1 + \sigma_{\pi a}^2 \varepsilon_{\pi a}) a + \theta_3^* \varepsilon_\pi \end{aligned}$$

where $\varepsilon_{\pi a}$ is an independent error and $\sigma_{\pi a}^2 > 0$. If we let the support of ε_U and ε_π be unbounded, nothing in the analysis below would be affected by a shift to this alternative specification. However, that change would bring us outside the compact Y setting study in the rest of the paper, so we opted for preserving the consistency.

towards cycles, a channel pointed out by Nyarko (1991) in a monopoly pricing setting.

Corollary 1. *There is $\bar{c} > 0$ such that for all $c \leq \bar{c}$*

1. *There is no c -robust equilibrium.*
2. *There exists a mixed c -robust equilibrium.*
3. *The maximal and minimal equilibria are such that $\alpha^*(0)$ is increasing in $\theta_{1\pi}^* + \theta_{1a}^*$.*

Playing the conservative policy is the best reply to a high misspecification concern and θ^* but induces a low concern as its consequences are well explained. In contrast, the aggressive policy is a best reply to a low misspecification concern and θ^* but induces a severe concern. Therefore, the policy cannot stabilize, consistently with the cyclical behavior of monetary policies documented in Sargent (1999), Clarida, Gali, and Gertler (2000), and Sargent (2008). We also have some natural comparative statics in the extremal robust equilibria, as a more significant exploitable trade-off between inflation and unemployment induces more time spent using an aggressive monetary policy.²⁴

In this application, we purposefully chose one of the most straightforward macroeconomic frameworks to isolate and illustrate the effect of an endogenous concern for misspecification. However, incorporating an endogenous concern for misspecification in more elaborate models is a valuable enterprise. For example, the fact that evidence impacts the trust in the model may be used to explain the observed pattern of initial underreaction to information when only beliefs within a model are adjusted, and medium-run overreaction, when the belief adjustment compounds with a change in model trust (see Angeletos, Huo, and Sastry, 2021 and the references therein for a discussion of this pattern).

5 Representation

We next move to characterize the average robust control model in terms of observable choices in an Anscombe-Aumann framework. In line with the literature on decision theory under uncertainty, our goal is to associate the decision criterion in equation (1) with axioms on a binary preference relation over acts.

²⁴It is well-known that non-extremal equilibria are less well-behaved in terms of comparative statics. See Diamond (1982) for a very early example. In any case, a supermodularity condition between the concern for misspecification and the conservative policy payoff guarantees equilibrium uniqueness.

Before jumping into the details of the axiomatization, we provide a high-level description of the steps involved and the intuitive meaning of the axioms we link to the representation. In terms of observability requirements, we allow the analyst to elicit preferences for bets both on the data-generating process, e.g., the urn composition, and on the actual realization, e.g., the color of the drawn ball.²⁵ The analysis then has three nested levels: 1) An axiomatization of the static decision criterion, 2) An axiomatization of the qualitative changes of the preference parameters, 3) An axiomatization of the speed of adjustment of the concern for misspecification (i.e., for the different β s).

The static decision criterion belongs to the variational class of Maccheroni, Marinacci, and Rustichini (2006a). More importantly, within this class, it is identified by a relaxed Sure-Thing Principle: the agent satisfies it for bets that involve the identity of the model (e.g., bets on an urn composition) and for bets on events conditional on the model (e.g., bets on the color after having revealed the urn composition). However, failures of the Sure-Thing principle can realize for acts that involve the realization of the outcome without conditioning on the model (e.g., bets on the color without knowing the urn composition, which are the ones involved in the classical Ellsberg’s paradox). The final conceptual axiom involved in the representation of equation (1) is a notion of uniform conditional misspecification concern. It requires that conditional on being told the identity of their best-fitting model, the agent is equally concerned about it not being exact regardless of which one it is.

We consider a collection of binary relations indexed by the observed history to characterize the agent’s dynamic preferences. Three other axioms identify the qualitative changes of the preference parameters u, λ, μ . Constant Preference Invariance guarantees that the taste u for uncertain alternatives is stable over time. Dynamic Consistency over Models guarantees that the probability distribution over models is updated in a Bayesian fashion. In contrast, Q -Likelihood guarantees that the concern for misspecification is increasing in the value of the likelihood ratio statistic within each time horizon.

Finally, we axiomatize the asymptotic speed of adjustment of the misspecification concern. To do so, we need a quantitative notion of how similar two preference rela-

²⁵This is standard when dealing with multiple sources of uncertainty, see for example Klibanoff, Marinacci, and Mukerji (2005) for an early example, and Cerreia-Vioglio, Maccheroni, Marinacci, and Montrucchio (2013a) for a general framework and results. We discuss how to relax this requirement in Section 6.5.

tions are, which is defined using an event E and two deterministic and strictly ranked outcomes, x and y , as measuring rods. Loosely speaking, two relations are (x, y, E, ε) similar if their certain equivalents for the binary act xEy are ε close. With this, an Asymptotic Frequentism axiom singles out the statistically sophisticated type: for every (x, y, E, ε) , the conditional preferences after sufficiently long sequences of outcomes sharing the same empirical frequency must be (x, y, E, ε) -similar. Conversely, a lenient type asymptotically becomes similar to those SEU preferences that are less misspecification concerned than the initial preference. The demanding type must approach the preferences of a maxmin agent, thus confirming in a decision-theoretic setting the insights of Theorem 1.

5.1 Notation and Preliminaries

The agent evaluates simple acts, i.e., measurable and finite ranged maps from a nonempty state space S into a convex set of outcomes X , where S is endowed with a σ -algebra of events Σ . The set of those acts is denoted as \mathcal{F} . Given any $x \in X$, $x \in \mathcal{F}$ is the act that delivers x in every state, and in this way, we identify X as the subset of constant acts in \mathcal{F} . If $f, g \in \mathcal{F}$, and $E \in \Sigma$, we denote as gEf the simple act that yields $g(s)$ if $s \in E$ and $f(s)$ if $s \notin E$. Since X is convex, for every $f, g \in \mathcal{F}$, and $\gamma \in (0, 1)$, we denote as $\gamma f + (1 - \gamma)g \in \mathcal{F}$ the simple act that pays $\gamma f(s) + (1 - \gamma)g(s)$ for all $s \in S$.

We model the agent's preference with a binary relation \succsim on \mathcal{F} . As usual \succ and \sim denote the asymmetric and symmetric parts of \succsim . An event E is *null* if $fEh \sim gEh$ for every $f, g, h \in \mathcal{F}$. An event is *nonnull* if it is not null. For every $E \in \Sigma$, the conditional preference relation \succsim_E is defined by $f \succsim_E g$ if $fEh \succsim gEh$ for some $h \in \mathcal{F}$.

A key concept to understand the concern for misspecification evolution is a notion of being more misspecification concerned from Ghirardato and Marinacci (2002).

Definition 8. Given two preferences \succsim_1 and \succsim_2 on \mathcal{F} , we say that \succsim_1 is more concerned with misspecification than \succsim_2 if, for each $f \in \mathcal{F}$ and each $x \in X$, $f \succsim_1 x$ implies $f \succsim_2 x$.

5.2 Decision Criterion

When formalized in terms of a binary relation, the average robust control decision criterion reads as follows.

Definition 9. A tuple (u, Q, μ, λ) is an *average robust control representation* of the preference relation \succsim if $u : X \rightarrow \mathbb{R}$ is a nonconstant affine function, $Q \subseteq \Delta(S)$ is a nonempty set, $\mu \in \Delta(Q)$, $\lambda \geq 0$, and for all $f, g \in \mathcal{F}$

$$f \succsim g \iff \mathbb{E}_\mu \left[\min_{p \in \Delta(S)} \int_S u(f) dp + \frac{1}{\lambda} R(p||q) \right] \geq \mathbb{E}_\mu \left[\min_{p \in \Delta(S)} \int_S u(g) dp + \frac{1}{\lambda} R(p||q) \right]. \quad (4)$$

The average robust control representation is the counterpart of (1) when expressed over acts. An apparent difference is that u here takes as input only outcomes instead of pair of actions and consequences. However, this discrepancy is inconsequential, as in Section 2 we can define a larger space of consequences $\hat{Y} = A \times Y$ that includes both actions and outcomes and transforming each model $p \in \Delta(Y)^A$ into an element of $\hat{p} \in \Delta(\hat{Y})^A$ such that $\hat{p}_a(a', y) = 0$ if $a' \neq a$ and $\hat{p}_a(a, y) = p_a(y)$ for all $y \in Y$. Still, this embedding of actions into outcomes muddles the interpretation of the learning results significantly. Therefore we opted to maintain the distinction explicit at the cost of some visual discrepancy between equations (1) and (4).²⁶

5.3 Static Axioms

Our first axiomatic step is a static one. We characterize in terms of behavioral axioms an agent that evaluates accordingly to equation (4) the acts whose consequences are obtained in the same period and before any new information is received.

Axiom 1 (Variational Axiom). **Weak Order.**

Weak Certainty Independence. If $f, g \in \mathcal{F}$, $x, x' \in X$, $\gamma \in (0, 1)$, and $\gamma f + (1 - \gamma)x \succsim \gamma g + (1 - \gamma)x$, then $\gamma f + (1 - \gamma)x' \succsim \gamma g + (1 - \gamma)x'$.

Continuity. If $f, g, h \in \mathcal{F}$ the sets $\{\gamma \in [0, 1] : \gamma f + (1 - \gamma)g \succsim h\}$ and $\{\gamma \in [0, 1] : h \succsim \gamma f + (1 - \gamma)g\}$ are closed.

Monotonicity. If $f, g \in \mathcal{F}$, and $f(s) \succsim g(s)$ for all $s \in S$, then $f \succsim g$.

Uncertainty Aversion. If $f, g \in \mathcal{F}$, $\gamma \in (0, 1)$, and $f \sim g$, then $g + \gamma(f - g) \succsim f$.

Nondegeneracy. $f \succ g$ for some $f, g \in \mathcal{F}$.

Weak Monotone Continuity. If $f, g \in \mathcal{F}$, $x \in X$, $(E_n)_{n \in \mathbb{N}} \in \Sigma^{\mathbb{N}}$ with $f \succ g$, $E_1 \supseteq E_2 \supseteq \dots$ and $\cap_{n \in \mathbb{N}} E_n = \emptyset$, then there exists $n_0 \in \mathbb{N}$ such that $xE_{n_0}f \succ g$.

²⁶See Fishburn (1970) Chapter 12.1 for a more detailed discussion of the equivalence of a formulation with exogenously given states and one where states are maps from actions into consequences.

Maccheroni, Marinacci, and Rustichini (2006a) shows that Axiom 1 characterizes the class of variational preferences. Weak Order, Continuity, and Nondegeneracy are standard technical requirements. Weak Monotone Continuity guarantees that the probabilistic scenarios considered by the agent are countably additive. Weak Certainty Independence allows the agent to perceive some advantage in hedging, but this cannot come from mixing with different constants using the same weights. Monotonicity requires that the preference over acts is minimally consistent with the preference over the outcomes they induce. Uncertainty Aversion leads to aversion for the acts that perform well for a postulated model but poorly for its perturbations.

5.3.1 Structured Preferences

We are considering agents who face two levels of uncertainty: the uncertainty on the best structured description of the data-generating process and whether each description is exact. A representation is structured if it allows separating these two layers. In particular, to achieve this separation, we consider a state space S that admits the decomposition $S = \Omega \times \Delta(\Omega)$ for some finite Ω endowed with its Borel sigma-algebra.

Definition 10. An average robust control representation (u, Q, μ, λ) is *structured* if μ has finite support and there exists a continuous injective map

$$\begin{aligned} Q &\rightarrow \Delta(\Omega) \\ q &\mapsto \rho_q \end{aligned}$$

such that for every $q \in Q$ and $\omega \in \Omega$, $q(\{\omega, \rho_q\}) = \rho_q(\omega)$.

The interpretation of a structured representation is that the state space can be factored in two components, the realization of the single period consequence $\omega \in \Omega$ and a component $\rho \in \Delta(\Omega)$ that pins down the distribution over states each period. An event E is *structured* if $E = \Omega \times B$ for some $B \in \mathcal{B}(\Delta(\Omega))$. The sigma-algebra generated by the structured events is denoted as Σ_s .²⁷

We say that an event $E \subseteq S$ satisfies the sure-thing principle if, for all $f, g, h, h' \in \mathcal{F}$ we have that $fEh \succsim gEh$ implies $fEh' \succsim gEh'$. We denote by Σ_{st} the set of events that satisfy the sure-thing principle.

²⁷With a slight abuse of notation for every $B \in \mathcal{B}(\Delta(\Omega))$ and $W \subseteq \Omega$ we denote as \succsim_B and \succsim_W the binary relations $\succsim_{\Omega \times B}$ and $\succsim_{W \times \Delta(\Omega)}$ and we write fBg and fWg for $f(\Omega \times B)g$ and $f(W \times \Delta(\Omega))g$.

Axiom 2 (Structured Savage). *i) There is a finite set $E \subseteq S$ such that $S \setminus E$ is null. ii) **P2**. $\Sigma_s \subseteq \Sigma_{st}$. iii) **P4**. If $E, E' \in \Sigma_s$ and $x, y, w, z \in X$ are such that $x \succ y$ and $w \succ z$, then*

$$xEy \succ xE'y \Rightarrow wEz \succ wE'z.$$

Structured Savage requires that (i) the agent posits a finite number of models and (ii) guarantees that when evaluating acts that only depend on the identity of the structured model, the agent satisfies the Sure-Thing Principle.²⁸ It also (iii) guarantees that when an agent faces alternatives whose outcomes depend only on whether the DGP belongs to two sets of models, their choices consistently reveal the one deemed more likely.

Axiom 3 (Intramodel Sure-Thing Principle). *For every $f, g, h, h' \in \mathcal{F}$,*

$$fWh \succsim_\rho gWh \implies fWh' \succsim_\rho gWh' \quad \forall W \subseteq \Omega, \forall \rho \in \Delta(\Omega).$$

Structured Savage's P2 and the Intramodel STP imply that bets *between* models and preference over acts *within* a model satisfy the STP. However, they admit violations of the STP for acts whose payoff depends on both the model's identity and the outcome realization within the model, as the ones of the original Ellsberg's paradox.

The case we study is when the relative likelihood of the structured models is only captured by the belief μ . In particular, the agent is equally concerned about how much each model departs from the actual data-generating process.

Axiom 4 (Uniform Misspecification Concern). *For every $\rho, \rho' \in \Delta(\Omega)$ and $f, g \in \mathcal{F}$ such that*

$$\rho(\{\omega : f(\omega, \rho) = y\}) = \rho'(\{\omega : g(\omega, \rho') = y\}) \quad \forall y \in X$$

and $\Omega \times \{\rho\}, \Omega \times \{\rho'\}$ are nonnull we have

$$f \succsim_\rho x \iff g \succsim_{\rho'} x \quad \forall x \in X.$$

This axiom requires that if acts f and g induce identical outcome distributions under ρ and ρ' , they are compared with a safe alternative in the same way conditional on the best fitting model being revealed to be ρ or ρ' .

²⁸The extension to infinitely many models does not provide additional conceptual difficulties but makes the conditioning involved in the dynamic axioms much more cumbersome. Gul and Pesendorfer (2014) introduces the idea of sources of uncertainty for which the decision maker can quantify uncertainty and connects it with the Sure-Thing Principle.

Definition 11. The state space is Q -adequate if for every $x, y \in X, \omega, \omega' \in \Omega$, and $\rho \in \Delta(\Omega)$: (i) there exist $k \in (0, 1)$ and $(W_q)_{q \in Q} \in (2^\Omega)^Q$ such that $\rho_q(W_q) = k$, (ii) if $\{\omega\} \times \{\rho\}$ and $\{\omega'\} \times \{\rho\}$ are nonnull, then $x_{\{\omega\} \times \{\rho\}} y \sim x_{\{\omega'\} \times \{\rho\}} y$.

All the agent's structured models have an event with the same probability and are uniform over a model-specific set of outcomes. It is well-known that equal probability requirements are essential for probabilistic sophistication with respect to a finite measure over states to have a bite (see, e.g., Chew and Sagi, 2006). They can be relaxed if we allow for a continuum of states and models. The only role of (i) for us is to obtain a concern for misspecification λ that is not model dependent (i.e., not to have $(\lambda_q)_{q \in Q}$ in the representation) from Uniform Misspecification Concern, an axiom with an unmistakable flavor of probabilistic sophistication.

Axiom 5 (Uncertainty Neutrality Over Models). *Let $x, y, w, z \in X$, $\rho \in \Delta(\Omega)$, and $\gamma, \zeta \in (0, 1)$ with $\gamma + \zeta \in (0, 1)$. Then $[\gamma x + (1 - \gamma) y]_\rho w \sim y_\rho z$ if and only if $[(\gamma + \zeta) x + (1 - \gamma - \zeta) y]_\rho w \sim [\zeta x + (1 - \zeta) y]_\rho z$.*

Uncertainty Neutrality over Models guarantees that at the level of bets over models, the agent is “risk-neutral”, as changing the performance under ρ by $(x - y)\gamma$ has an impact that does not depend on the level of utility under that model. It is immediate from the proof of Theorem 3 that if dropped, it leads to a more general representation with a nonlinear utility index U over the performance of each robust control model.

Theorem 3. *Suppose that S is Q -adequate, there at least three disjoint nonnull events in Σ_s , and every nonnull $E \in \Sigma_s$ contains at least three disjoint nonnull events. The following are equivalent:*

1. \succsim admits a structured average robust control representation (u, Q, μ, λ) ;
2. \succsim satisfies Variational Axiom, Structured Savage, Uniform Misspecification Concern, Intramodel Sure-Thing Principle, and Uncertainty Neutrality over Models.

Moreover, in this case, every two structured average robust control representations share the same μ .

The theorem characterizes the representation (u, Q, μ, λ) with probabilistic uncertainty *about* the model (Structured Savage), probabilistic sophistication *given* a model (Intramodel Sure-Thing Principle), and *incomplete trust* in any model (Uncertainty Aversion).

Corollary 2. *Suppose that \succsim admits a structured average robust control representation (u, Q, μ, λ) . Then \succsim is more misspecification concerned than the subjective utility preference with utility index u and belief $\int_Q q d\mu(q)$.*

5.4 Dynamic Axioms

We next provide axioms that characterize the dynamic adjustment of preferences in the face of information. In particular, we look at joint axioms on a collection of history-dependent binary relations $(\succsim^h)_{h \in \mathcal{H}}$ indexed by the realized history. Recall that the relevant set of length $t \in \mathbb{N}$ histories for structured preferences is Ω^t .

Axiom 6 (Constant Preference Invariance). *For every $x, x' \in X$ and $h \in \mathcal{H}$,*

$$x \succsim^h x' \Leftrightarrow x \succsim^\emptyset x'.$$

This axiom captures the fact that we are not considering the problem of an agent discovering their taste. The preferences over uncertain alternatives are fixed and do not react to new information.

Axiom 7 (Dynamic Consistency over Models). *Let $f, g \in \mathcal{F}$ be Σ_s -measurable, $t \in \mathbb{N}$, $(\omega_1, \dots, \omega_t) \in \Omega^t$ and $\bar{z}, \underline{z} \in X$ be such that $\bar{z} \succsim f(s) \succsim \underline{z}$ and $\bar{z} \succsim g(s) \succsim \underline{z}$ for all $s \in S$. Define h^0 as*

$$h^0(\omega, \rho) = \gamma_{h(\omega, \rho)} \prod_{i=1}^t \rho(\omega_i) \bar{z} + \left(1 - \gamma_{h(\omega, \rho)} \prod_{i=1}^t \rho(\omega_i)\right) \underline{z} \quad \forall (\omega, \rho) \in S, \forall h \in \{f, g\}$$

where $\gamma_{h(\omega, \rho)}$ satisfies $h(\omega, \rho) = \bar{z} \gamma_{h(\omega, \rho)} + (1 - \gamma_{h(\omega, \rho)}) \underline{z}$. Then, we have

$$f \succsim^{(\omega_1, \dots, \omega_t)} g \iff f^0 \succsim g^0.$$

The second dynamic axiom requires Bayesian rationality when considering acts whose consequences only depend on the structured model. Formally, it requires that when comparing acts that only bet on the identity of the model, at a given history, we can reduce the comparison to acts evaluated ex-ante. To do so, the payoff conditional to each model must be scaled proportionally to the amount of evidence that has been generated in favor of that model.²⁹ To link the misspecification concern to the evidence,

²⁹This axiom can lead to fruitful implications beyond our average robust control decision criterion,

we need to translate the notion of the likelihood ratio test for a set of models in our setting.³⁰

$$LRT(h_t, Q) = -2 \log \left(\frac{\max_{q \in Q} \prod_{\tau=1}^t \rho_q(\omega_\tau)}{\max_{\rho \in \Delta(\Omega)} \prod_{\tau=1}^t \rho(\omega_\tau)} \right).$$

Axiom 8 (Q -Likelihood). *For every $t \in \mathbb{N}$, $h_t, h'_t \in \mathcal{H}_t$, and $\rho \in \Delta(\Omega)$ if*

$$LRT(h_t, Q) \geq LRT(h'_t, Q)$$

then $\succsim_\rho^{h_t}$ is more concerned with misspecification than $\succsim_\rho^{h'_t}$.

The third dynamic axiom links the fit of the models in Q with the evolution of the misspecification concern *within* each period t . In particular, it requires that the concern for misspecification increases in the likelihood ratio test statistic. Crucially, the comparison must be performed at the level of the conditional preferences, as different posteriors over structured models imply that the unconditional preferences are generally not rankable in terms of more misspecification aversion.³¹

Proposition 4. *Let $(\succsim^h)_{h \in \mathcal{H}}$ and $Q \subseteq \Delta(S)$ be such that:*

1. *For every $h \in \mathcal{H}$, \succsim^h satisfies the axioms of Theorem 3,*
2. *$(\succsim^h)_{h \in \mathcal{H}}$ satisfies Constant Preference Invariance, Q -Likelihood, and Dynamic Consistency over Models.*

Then for every $h \in \mathcal{H}$, \succsim^h admits an average robust control representation $(u, Q, \mu(\cdot|h), \lambda_h)$ with λ_h increasing in $LRT(\cdot, Q)$ among histories of the same length.

The comparative notion of more misspecification averse employed in the Q -Likelihood axiom is ordinal. Therefore, as witnessed by the previous proposition, it allows ordering the concern for misspecification for different realizations of the likelihood test. However, to single out the *quantitative* speed at which the concern for misspecification is adjusted, we need a quantitative measure of similarity.

as it implies Bayesian updating for each decision criteria that performs an average of model-specific evaluations (that could, for example, take the form of other divergence preferences or rank-dependent utility evaluations). In this way, it would complement the elegant theory of subjective learning developed in Dillenberger, Lleras, Sadowski, and Takeoka (2014), which does not require that the analyst observes the same information as the agent.

³⁰Given the finiteness of Ω , we can focus on the case in which the alternative set of models $N(Q) = \Delta(\Omega)$, discussed in Section 2. The extension to general sets of alternative models is straightforward.

³¹A somewhat related idea of comparing attitudes conditional on some information that is perceived as objective is also present in Frick, Iijima, and Le Yaouanq (2022).

For every $x, y \in X$ with $x \succ y$ and $E \in \Sigma$ let γ_{\succ}^{xEy} be defined by

$$\gamma_{\succ}^{xEy} x + (1 - \gamma_{\succ}^{xEy}) y \sim xEy.$$

That is, γ_{\succ}^{xEy} is the weight to alternative x in the certain equivalent to act xEy . It captures both the confidence in event E and the attitudes towards uncertainty. It is easy to see that under the Variational Axiom γ_{\succ}^{xEy} always exists and is unique.

For every $x, y \in X$, $E \in \Sigma$, $\varepsilon \in (0, 1)$, and \succ and \succ' that satisfy the Variational Axiom, we say that \succ is (x, y, E, ε) -similar to \succ' if

$$\left| \gamma_{\succ}^{xEy} - \gamma_{\succ'}^{xEy} \right| \leq \varepsilon.$$

That is, the certain equivalent of the binary act xEy is ε close under preferences \succ and \succ' .

Axiom 9 (Asymptotic Frequentism). *For every $\rho \in \Delta(\Omega)$, $x, y \in X$ with $x \succ^{\emptyset} y$, $\varepsilon \in (0, 1)$, and $E \in \Sigma$ there is $\tau \in \mathbb{N}$ such that if $t, t' \geq \tau$ and $h_t, h_{t'}$ have outcome frequency ρ then \succ^{h_t} is (x, y, E, ε) -similar to $\succ^{h_{t'}}$.*

The axiom requires that for every binary act xEy , a sufficiently long sequence of outcomes with the same empirical frequency stabilize the certain equivalent.

Proposition 5. *If $(\succ^h)_{h \in \mathcal{H}}$ satisfies the assumptions of Proposition 4 and Asymptotic Frequentism then for every sequence $(h_{t_n})_{n \in \mathbb{N}}$ with a constant outcome frequency not in $\{\rho_q : q \in Q\}$,*

$$\lim_{n \rightarrow \infty} \lambda_{h_{t_n}} / \left(\frac{LRT(h_{t_n}, Q)}{2t_n} \right) \quad (5)$$

exists. Moreover, if for some $q \in Q$ and nonnull $E \subseteq \Omega \times \{\rho_q\}$,

$$\liminf_{n \rightarrow \infty} \gamma_{\succ^{h_{t_n}}}^{xEy} > 0,$$

the limit is finite, and if

$$\limsup_{n \rightarrow \infty} \gamma_{\succ^{h_{t_n}}}^{xEy} < \rho_q(E) \mu(q),$$

it is strictly positive.

To better grasp this result, recall that in our learning model $\lambda_{h_{t_n}} = \frac{LRT(h_{t_n}, Q)}{2\beta_{t_n}}$. Therefore requiring that the limit in equation (5) equals some constant $1/c$ implies that

$\lim_t \frac{\beta_t}{t} = c$, i.e., it pins down an agent who asymptotically adjusts its misspecification concern as a statistically sophisticated type.

The proof of the result has two main steps. First, we show that the likelihood ratio statistic of the models Q is growing linearly in t_n along the sequence of histories $(h_{t_n})_{n \in \mathbb{N}}$, so that the denominator in equation (5) converges. Because the outcome frequency does not correspond to a model in Q , this limit is not 0. With this, the proof amounts to showing that the revealed concern for misspecification also converges. The second step rules out the existence of different finite limit points for $(\lambda_{t_n})_{n \in \mathbb{N}}$ by contradiction. If these points exist, then for every pair of strictly ranked outcomes $x \succ y$, we construct an event E for which the DM does not satisfy the Sure-Thing Principle such that the preference with the high concern for misspecification has a strictly higher certain equivalent than the one with the low concern.

Axiom 10 (Asymptotic Concern). *Let $f \in \mathcal{F}$, $x \in X$, and $\hat{\rho}, \rho \in \Delta(\Omega)$ be such that $\Omega \times \{\hat{\rho}\}$ is \succsim^\emptyset -null, $\Omega \times \{\rho\}$ is \succsim^\emptyset -nonnull, and $\rho \gg \hat{\rho}$. If $\rho(\{\omega \in \Omega : x \succ f(\omega, \rho)\}) > 0$, then there exists $\tau \in \mathbb{N}$ such that for all $t \geq \tau$ and all h_t with outcome frequency $\hat{\rho}$, $x \succ_{\rho}^{h_t} f$.*

Asymptotic Concern requires that long-run failures in explaining the data (i.e., an empirical frequency $\hat{\rho}$ that is not among the agent's structured models) increase the concern so that every certain outcome is preferable to an act with worse payoffs under a relevant model ρ .

Proposition 6. *If $(\succsim^h)_{h \in \mathcal{H}}$ satisfies the assumptions of Proposition 4 and Asymptotic Concern then for every sequence $(h_{t_n})_{n \in \mathbb{N}}$ with a constant outcome frequency that is not in $\{\rho_q : q \in Q\}$ we have*

$$\frac{LRT(h_{t_n}, Q)}{\lambda_{h_{t_n}}} = o(t_n).$$

This result shows that Asymptotic Concern characterizes agents who apply an excessively demanding time discount to the likelihood ratio test statistic (see equation (2)). Indeed, the time scaling elicited as the ratio between the LRT and the concern for misspecification revealed by the choices grows sublinearly time, the condition that defines demanding agents.

Axiom 11 (Asymptotic Leniency). *Let $x, y \in X$, $E \in \Sigma$, $\rho \in \Delta(\Omega)$, $\varepsilon \in (0, 1)$, $B \subseteq \Delta(\Omega)$, be such that $x \succ y$ and $S \setminus E, \Omega \times B$ are nonnull. For every Bayesian SEU*

preferences $(\succeq^h)_{h \in \mathcal{H}}$ such that \succeq_B^\emptyset is less misspecification averse than \succsim_B^\emptyset , there exists $\tau \in \mathbb{N}$ such that for every $t \geq \tau$ and h_t with outcome frequency ρ , \succsim^{h_t} is (x, y, E, ε) -similar to \succeq^{h_t} .

Asymptotic Leniency requires that if the empirical distribution converges to some $\rho \in \Delta(\Omega)$, the preferences of the agents approximate, i.e., are eventually (x, y, E, ε) -similar to the updated preferences of an SEU whose model contingent preferences were initially less misspecification averse than the agent.

Proposition 7. *If $(\succsim^h)_{h \in \mathcal{H}}$ satisfies the assumptions of Proposition 4 and Asymptotic Leniency then for every sequence $(h_{t_n})_{n \in \mathbb{N}}$ with constant outcome frequency not in $\{\rho_q : q \in Q\}$*

$$t_n = o\left(\frac{LRT(h_{t_n}, Q)}{\lambda_{h_{t_n}}}\right).$$

This proposition shows that convergence to subjective expected utility maximization (in the form of Asymptotic Leniency) characterizes excessively lenient time normalizations.

6 Discussion

6.1 Related Literature

A few papers allow the agents to realize that they are misspecified. In particular, in He and Libgober (2022), Ba (2022), Fudenberg and Lanzani (2022), and Gagnon-Bartsch, Rabin, and Schwartzstein (2022) misspecification can be eliminated either by “light bulb realizations” or evolutionary pressure. The key difference with our approach is that in these papers, as well in the earlier Foster and Young (2003), Cho and Kasa (2015), and Giacomini, Skreta, and Turén (2015), where agents switch between models on the basis of a specification test, the agents act as if they have complete trust in the set of models currently entertained and are never concerned about being misspecified. Still, there is a tight connection between the robust control decision criterion and a maxmin decision criterion where the set of models expands as the penalization term in the robust control increases (see Hansen and Sargent, 2011, for a textbook treatment). In light of this, compared to the previous set of papers, our work can additionally be interpreted as providing the first *smooth* framework for expanding (or restricting) the set of considered models as a function of the evidence. Farther afield, Ortoleva (2012)

proposes and axiomatizes a model where a decision maker can reject their model in favor of a backup one when faced with events with sufficiently low probability. Karni and Vierø (2013) proposes and axiomatizes a model where the agent becomes progressively aware of more states and acts. However, their decision maker trusts their probability over states completely when making decisions. Banerjee, Chassang, Montero, and Snowberg (2020) studies a Wald problem where the agent trade-offs between robustness and the subjective expected utility performance of the experiment. Differently from us, the concern in this model does not evolve, and the agent makes a single decision. Epstein and Ji (2022) characterizes optimal stopping with a concern for robustness captured by maxmin preferences, showing that the robustness concern, in general, induces earlier stopping.

There is fast-growing literature on learning under misspecification with subjective expected utility preferences. Arrow and Green (1973) gives the first general framework for this problem, and Nyarko (1991) points out that the combination of misspecification and endogenous data can lead to cycles. This literature has been revived by the more recent Esponda and Pouzo (2016); see Bohren and Hauser (2021), Esponda, Pouzo, and Yamamoto (2021a), Fudenberg, Lanzani, and Strack (2021), and Frick, Iijima, and Ishii (2022) for analyses of more closely related settings.

The identification of an agent who is disappointed with minor discrepancies between the empirical and the theoretical distributions as a believer in the Law of Small Numbers follows the formalization of this bias proposed by Rabin (2002). The normative role of a scaling of the likelihood ratio that makes it proportional to the relative entropy (cf. Proposition 1 and Theorem 1) is somewhat reminiscent of the normative role of (absolute) entropy as a measure of informativeness found by Cabrales, Gossner, and Serrano (2013).

Hansen and Sargent (2007) mention a time-varying penalization parameter as a way to maintain dynamic consistency in the robust control model. Maenhout (2004) also uses a time-varying penalization parameter in a portfolio selection problem to keep the recursive discounted preferences homothetic at any history. In both cases, the parameter evolution does not capture the fit of the models to the observed data. In the evolutionary game theory literature studying preferences formation, the closer papers are Dekel, Ely, and Yilankaya (2007) and Robatto and Szentes (2017). In the former, players may be “misspecified” in the sense of having a vN-M utility different from the one determining reproductive finiteness. In the latter, the evolutionary pressure determines the risk attitudes of the players.

On the axiomatic side, the static decision criterion considered here is due to Cerreia-Vioglio, Hansen, Maccheroni, and Marinacci (2022).³² The explicit use of a state space where every state describes both the single-period outcome realization and the probability distribution over outcomes follows the approach introduced in Cerreia-Vioglio, Maccheroni, Marinacci, and Montrucchio (2013a) as a two-stage “statistical” interpretation and axiomatization of some of the decision criteria under ambiguity, in particular the smooth ambiguity one. For this criterion, this approach has been recently extended by Denti and Pomatto (2022). They allow for a fully revealed-preference elicitation of the relevant probability distributions, viewed as subjective statistical models. See also Dean and Ortaleva (2017) for a less related decision criterion where the agent has a prior over multiple data-generating processes and evaluates each of them with rank-dependent utility and Gilboa, Minardi, and Samuelson (2020) for a different quasi-Bayesian criterion that combines Bayesian updating with case-reasoning rather than misspecification considerations.

6.2 Experimental Evidence

We are unaware of experiments that explicitly test the positive relation between misspecification and the belief in the Law of Small Numbers with uncertainty aversion. However, the findings in Esponda, Vespa, and Yuksel (2022) suggest that a mechanism similar to the one outlined in this paper is actually at play. The paper studies the repeated behavior of two groups of agents, one with an agnostic (full support) belief about the possible data-generating process faced and one that is misspecified because of base rate neglect. The average play of the misspecified agents is in between the best reply to the misspecified model and the uniform distribution over outcomes. Notably, this behavior is not the best reply to the observed empirical frequency, which suggests that, as in our model, even in the medium run (200 repetitions in the experiment above), the agents do not altogether drop their models; they rely less on it to make their choices. Instead, the correctly specified agents converge to making choices that are optimal only under the actual data-generating process, i.e., they behave as a subjective expected utility maximizer with a belief concentrated on the true DGP.

³²Given the Donsker–Varadhan variational formula, our decision criterion can also be seen as the average of CARA certain equivalents, an object studied and characterized from a statistical perspective in Mu, Pomatto, Strack, and Tamuz (2021).

6.3 General Endogenous Concern

One natural extension of the model allows for a more general form of dependence of the concern for misspecification than the one assumed in equation (2) and relaxes the ratio structure. A first generalization has $\lambda(h_t) = F(t, LRT(h_t, \Theta))$ where $F : \mathbb{N} \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is a function that maps a period and a realization of the likelihood-ratio statistic into a level of misspecification concern. The main learning results extend, with the statistically sophisticated behavior pinned down by the condition that for each $k > 0$, $\lim_{n \rightarrow \infty} F(n, kn)$ exists, is finite, and is increasing in k , while demanding and lenient agents correspond to the cases in which these limits are always equal to infinity and 0, respectively.

A more substantial generalization would be to allow the F above to depend on a function of the likelihood ratio different than the LRT . By applying a log function to the likelihood ratio, the LRT penalizes equally failures in all periods. To capture the exact dual pattern of behavior of believers in the Law of Small Numbers, i.e., an excessively demanding behavior in the short run and a more permissive one in the long run, one may want to allow for a concave transformation of the log function. In that case, one of the more interesting aspects to explore is the difference in uncertainty attitudes between the medium and long run.

6.4 Forward-looking Agents

One key generalization to our model would be to allow for forward-looking agents. Of course, as for many decision criteria that depart from SEU, the main complication is dealing with the fact that the most immediate extension of the criterion to forward-looking agents would induce dynamic inconsistencies under some information structures. One approach would be to directly impose a recursive formulation for the preferences, as in Maccheroni, Marinacci, and Rustichini (2006b) and Klibanoff, Marinacci, and Mukerji (2009). Since the decision criterion belongs to the variational class, we know from the first reference that a recursive formulation can be obtained. A complementary approach does not impose recursivity and allows for dynamic inconsistency. Preliminary analysis suggests that if we endow the agent with commitment power as in Hansen and Sargent (2005), or if we consider naive agents who do not anticipate their future taste variations, little is changed. However, analyzing an uncommitted, forward-looking, and sophisticated agent playing an intra-personal equilibrium with their future selves would require combining the insights of this paper with the approach developed

in Battigalli, Francetich, Lanzani, and Marinacci (2019). Analogously, to extend the axiomatic exercise to sophisticated agents, the techniques of this paper should be combined with the consistent planning approach of Siniscalchi (2011).

6.5 Endogenous Structured Models

The more natural extension for the decision-theoretic part of the paper involves using axioms that do not explicitly allow the agent to bet on the identity of the structured model. Allowing such bets is relatively standard when dealing with two levels of uncertainty for which the agent has different attitudes (Klibanoff, Marinacci, and Mukerji, 2005 being the most prominent example). However, Denti and Pomatto (2022) proposed an identifiability condition that avoids the need for explicit bets on the structured models. Identifiability requires a way to partition S that singles out the probabilistic model. In particular, each probabilistic model assigns probability one to its corresponding partition element.

When considering a structured environment, we required that this identification is spelled out in the description of the states, with the second component being the distribution ρ . Although in light of the results of Denti and Pomatto (2022), the axiomatization of this static criterion without this restriction does not generate conceptual complications, the dynamic characterization becomes significantly more involved. In particular, the challenge is created by the conditioning with respect to the endogenously identified model. This substantial extension is left for future work.

7 Conclusion

In this work, we propose a novel model of agents actively learning about the environment and dynamically adjusting their concern for misspecification based on the evidence they face. We show that the agents develop different long-run uncertainty attitudes depending on their understanding of how quickly evidence in favor or against a model is accumulated. Statistically sophisticated agents converge to robust control preferences a la Hansen and Sargent (2001), with the misspecification concern endogenously determined by their models fit with the true DGP at the equilibrium action. In contrast, an agent who is too demanding in evaluating their model converges to behave as a maxmin agent a la Gilboa and Schmeidler (1989), while a lenient agent eventually becomes a standard subjective expected utility maximizer. These results provide the

first learning foundation for nonstandard decision criteria.

We then point out that in natural environments, the behavior of the statistically sophisticated type need not converge, and we characterize the limit frequency of time spent playing each action. We apply this result to a simple macroeconomic model and obtain an explanation for the periodic switches in monetary policies.

We also provide an axiomatization of the proposed decision criterion and its evolution in the face of evidence. We leverage the static representation and a concept of preference similarity to introduce a new axiom type, Asymptotic Frequentism, requiring long streams of outcomes with the same empirical frequency to induce similar preferences. We prove that this axiom induces the statistically sophisticated behavior studied in the learning part of the paper.

A Appendix

A.1 Learning Results

A.1.1 Preliminaries

Let $p_a \in \Delta(Y)$ denote the a -th component of $p \in \Delta(Y)^A$, and

$$\|p - p'\| = \max_{a \in A} d(p_a, p'_a) \quad \forall p, p' \in \Delta(Y)^A$$

where d denotes the Prokhorov metric. By Assumption 1, there exists $K \in \mathbb{R}_{++}$ such that

$$-\ln \tilde{q}_a^\theta(y) \leq K \quad \forall a \in A, \forall y \in Y, \forall \theta \in \Theta.$$

Throughout Appendix A.1, the symbol K will denote such strictly positive real number.

For an arbitrary Borel measurable subset C of a metric space, we endow the space $C^{\mathbb{N}}$ with the Borel σ -algebra, $\mathcal{B}(C^{\mathbb{N}})$, corresponding to the product topology on $C^{\mathbb{N}}$. We denote by $k^t = (k_1, \dots, k_t)$ both the finite sequence in C^t and the elementary cylinder in $C^{\mathbb{N}}$ that it identifies. For every policy $\Pi \in A^{\mathcal{H}}$, the density of the objective probability distribution over infinite histories is defined over elementary cylinders as

$$\tilde{\mathbb{P}}_\Pi(a^t, y^t) = \begin{cases} 1 & \text{if } a_{\tau+1} = \Pi(a^\tau, y^\tau) \text{ for all } \tau \in \{0, \dots, t-1\}, \\ 0 & \text{otherwise,} \end{cases} \quad (6)$$

and $\mathbb{P}_\Pi(a^t, C) = \int_C \tilde{\mathbb{P}}_\Pi(a^t, \cdot) d(\prod_{\tau=1}^t p_{a_\tau}^*)$. Since the corresponding set of finite-dimensional

probability measures is consistent, there is a unique probability measure over infinite sequences of action-outcome pairs with these marginals, defined through the Kolmogorov extension theorem (see Theorem V.5.1 in Parthasarathy, 2005 for the version for general Borel spaces used here).

For every history $h_t = (a^t, y^t) \in \mathcal{H}$ let $p^{h_t} \in \Delta(Y)^A$ be the action contingent (finite support) probability measure over outcomes corresponding to the empirical frequency: for all $a \in A$ such that $\sum_{\tau=1}^t \mathbb{I}_{\{a\}}(a_\tau) > 0$,

$$p_a^{h_t}(C) = \frac{\sum_{\tau=1}^t \mathbb{I}_{\{(a,y): y \in C\}}(a_\tau, y_\tau)}{\sum_{\tau=1}^t \mathbb{I}_{\{a\}}(a_\tau)} \quad \forall C \subseteq \mathcal{B}(Y)$$

and $p_a^{h_t} = \delta_{\bar{y}}$ for some arbitrary fixed $\bar{y} \in Y$ if $\sum_{\tau=1}^t \mathbb{I}_{\{a\}}(a_\tau) = 0$. For every two histories $h_t, h_\tau \in \mathcal{H}$ we write $h_t \succ h_\tau$ if there is $n \in \mathbb{N}$ and $(a_i, y_i)_{i=1}^n$ such that $h_t = (h_\tau, (a_i, y_i)_{i=1}^n)$. For all $b \in A$ let Π^b the policy that prescribes b at every period.

A.1.2 Results

Our first lemma justifies the repeated use of min and max rather than inf and sup in the definitions of the average robust criterion and LRT.

Lemma 1. 1. For every $a \in A$, $\lambda \in \mathbb{R}_{++}$, and $\theta \in \Theta$

$$\emptyset \neq \operatorname{argmin}_{p_a \in \Delta(Y)} \mathbb{E}_{p_a}[u(a, y)] + \frac{1}{\lambda} R(p_a || q_a^\theta).$$

2. For every $t \in \mathbb{N}$, $h_t \in \mathcal{H}_t$,

$$\emptyset \neq \operatorname{argmax}_{p \in N(\Theta)} \prod_{\tau=1}^t \tilde{p}_{a_\tau}(y_\tau).$$

Proof. 1) Fix $a \in A$, $\lambda \in \mathbb{R}_{++}$, and $\theta \in \Theta$. Since u is continuous and Y is compact, $u(a, \cdot)$ is bounded, $\mathbb{E}_{p_a}[u(a, y)] \geq \min_{y \in Y} u(a, y) \in \mathbb{R}$ for all $p_a \in \Delta(Y)$, $p_a \mapsto \mathbb{E}_{p_a}[u(a, y)]$ is continuous, and since Y is a Polish space (being a compact metric space) $p_a \mapsto R(p_a || q_a^\theta)$ is lower semicontinuous by Lemma 1.4.3 in Dupuis and Ellis (2011). Therefore, the set

$$E := \left\{ p_a \in \Delta(Y) : \mathbb{E}_{p_a}[u(a, y)] + \frac{1}{\lambda} R(p_a || q_a^\theta) \leq \mathbb{E}_{q_a^\theta}[u(a, y)] \right\}$$

is closed, and as $R(q_a^\theta || q_a^\theta) = 0$ we clearly have

$$\operatorname{argmin}_{p_a \in \Delta(Y)} \mathbb{E}_{p_a} [u(a, y)] + \frac{1}{\lambda} R(p_a || q_a^\theta) = \operatorname{argmin}_{p_a \in \Delta(Y) \cap E} \mathbb{E}_{p_a} [u(a, y)] + \frac{1}{\lambda} R(p_a || q_a^\theta).$$

Since Y is compact, by Theorem 15.11 in Aliprantis and Border (2013) so are $\Delta(Y)$ and $\Delta(Y) \cap E$ endowed with the topology of weak convergence of measures. Since

$$p_a \mapsto \mathbb{E}_{p_a} [u(a, y)] + \frac{1}{\lambda} R(p_a || q_a^\theta)$$

is real-valued (with values in $[\min_{y \in Y} u(a, y), \mathbb{E}_{q_a^\theta} [u(a, y)]]$) and lower semicontinuous on the compact $\Delta(Y) \cap E$, it admits a minimizer by the generalized Weierstrass' theorem (see, e.g., Theorem 2.43 in Aliprantis and Border, 2013).

2) Let $t \in \mathbb{N}$, $h_t \in \mathcal{H}_t$. It follows from Assumption 2 (ii) and Theorem 1 in Sweeting (1986) that

$$\begin{aligned} N(\Theta) &\rightarrow \mathbb{R} \\ p &\mapsto \prod_{\tau=1}^t \tilde{p}_{a_\tau}(y_\tau) \end{aligned}$$

is continuous.³³ Then, the maximum is obtained since $N(\Theta)$ is closed by Assumption 2 (i), and thus compact by the compactness of Y . \blacksquare

The next lemma provides a useful rewriting of the updated concern for misspecification as a weighted average of the empirical likelihood ratio when playing the different actions, with weights proportional to how frequently each action has been used in the past.

Lemma 2. *For every $t \in \mathbb{N}$ and $h_t = (a^t, y^t) \in \mathcal{H}_t$, if $\hat{\theta} \in \operatorname{argmax}_{\theta \in \Theta} \prod_{\tau=1}^t \tilde{q}_{a_\tau}^\theta(y_\tau)$, and $p \in \operatorname{argmax}_{r \in N(\Theta)} \prod_{\tau=1}^t \tilde{r}_{a_\tau}(y_\tau)$ then*

$$LRT(h_t, \Theta) = \sum_{a \in A} \sum_{\tau=1}^t \mathbb{I}_{\{a\}}(a_\tau) \int_Y \log \left(\frac{dp_a(y)}{dq_a^{\hat{\theta}}(y)} \right) dp_a^{h_t}(y)$$

and

$$\lambda^\beta(h_t) = \sum_{a \in A} \sum_{\tau=1}^t \mathbb{I}_{\{a\}}(a_\tau) \int_Y \log \left(\frac{dp_a(y)}{dq_a^{\hat{\theta}}(y)} \right) dp_a^{h_t}(y) / \beta_t.$$

³³Observe that although the theorem in Sweeting (1986) is stated for densities with respect to the Lebesgue measure, both Scheffe's Theorem and Theorem 2.18 in Rudin (1970), that are used to prove the result, work for any regular Borel measure, as the $(p_a^*)_{a \in A}$ we consider are.

Proof. We have

$$\begin{aligned}
LRT(h_t, \Theta) &= -\log \left(\frac{\max_{\theta \in \Theta} \prod_{\tau=1}^t \tilde{q}_{a_\tau}^\theta(y_\tau)}{\max_{r \in N(\Theta)} \prod_{\tau=1}^t \tilde{r}_{a_\tau}(y_\tau)} \right) \\
&= \log \left(\frac{\max_{r \in N(\Theta)} \prod_{\tau=1}^t \tilde{r}_{a_\tau}(y_\tau)}{\max_{\theta \in \Theta} \prod_{\tau=1}^t \tilde{q}_{a_\tau}^\theta(y_\tau)} \right) = \log \left(\frac{\prod_{\tau=1}^t \tilde{p}_{a_\tau}(y_\tau)}{\prod_{\tau=1}^t \tilde{q}_{a_\tau}^{\hat{\theta}}(y_\tau)} \right) \\
&= \log \left(\prod_{\tau=1}^t \tilde{p}_{a_\tau}(y_\tau) \right) - \log \left(\prod_{\tau=1}^t \tilde{q}_{a_\tau}^{\hat{\theta}}(y_\tau) \right) \\
&= \log \left(\prod_{y \in Y} \prod_{a \in A} \tilde{p}_a(y)^{\sum_{\tau=1}^t \mathbb{I}_{\{a\}}(a_\tau) p_a^{h_t}(\{y\})} \right) \\
&\quad - \log \left(\prod_{y \in Y} \prod_{a \in A} \tilde{q}_a^{\hat{\theta}}(y)^{\sum_{\tau=1}^t \mathbb{I}_{\{a\}}(a_\tau) p_a^{h_t}(\{y\})} \right) \\
&= \sum_{a \in A} \sum_{\tau=1}^t \mathbb{I}_{\{a\}}(a_\tau) \sum_{y \in Y} p_a^{h_t}(\{y\}) \log(\tilde{p}_a(y)) \\
&\quad - \sum_{a \in A} \sum_{\tau=1}^t \mathbb{I}_{\{a\}}(a_\tau) \sum_{y \in Y} p_a^{h_t}(\{y\}) \log(\tilde{q}_a^{\hat{\theta}}(y)) \\
&= \sum_{a \in A} \sum_{\tau=1}^t \mathbb{I}_{\{a\}}(a_\tau) \left(\sum_{y \in Y} p_a^{h_t}(\{y\}) \log(\tilde{p}_a(y)) - \sum_{y \in Y} p_a^{h_t}(\{y\}) \log(\tilde{q}_a^{\hat{\theta}}(y)) \right) \\
&= \sum_{a \in A} \sum_{\tau=1}^t \mathbb{I}_{\{a\}}(a_\tau) \int_Y \log \left(\frac{\tilde{p}_a(y)}{\tilde{q}_a^{\hat{\theta}}(y)} \right) dp_a^{h_t}(y)
\end{aligned}$$

and the expression for $\lambda^{\mathbf{b}}(h_t)$ immediately follows by equation 2. ■

The next lemma shows that a robust control evaluation with respect to a structured model $\theta \in \Theta$ converges to a subjective expected utility evaluation as λ tends to 0, generalizing previous results in the decision-theoretic literature, where the function evaluated was a finite range one, to continuous utility functions.

Lemma 3. *For every $a \in A$ and $(\theta_n, \lambda_n)_{n \in \mathbb{N}} \in (\Theta \times \mathbb{R}_{++})^{\mathbb{N}}$ with*

$$\lim_{n \rightarrow \infty} (\theta_n, \lambda_n) = (\theta, 0)$$

we have

$$\lim_{n \rightarrow \infty} \min_{p_a \in \Delta(Y)} \mathbb{E}_{p_a} [u(a, y)] + \frac{1}{\lambda_n} R(p_a || q_a^{\theta_n}) = \mathbb{E}_{q_a^\theta} [u(a, y)].$$

Proof. Fix $a \in A$ and define $\bar{u} = \max_{y \in Y} u(a, y) - \min_{y \in Y} u(a, y)$. For every $n \in \mathbb{N}$,

$$\min_{p_a \in \Delta(Y)} \mathbb{E}_{p_a} [u(a, y)] + \frac{1}{\lambda_n} R(p_a || q_a^{\theta_n}) \in \left[\min_{y \in Y} u(a, y), \mathbb{E}_{q_a^\theta} [u(a, y)] \right],$$

so possibly restricting to a subsequence, we can assume that the limit in the LHS of the statement is well defined. The statement is then proved by showing that any such subsequence converges to the RHS. In particular, we show that we cannot have

$$\lim_{n \rightarrow \infty} \min_{p_a \in \Delta(Y)} \mathbb{E}_{p_a} [u(a, y)] + \frac{1}{\lambda_n} R(p_a || q_a^{\theta_n}) < \mathbb{E}_{q_a^\theta} [u(a, y)].$$

If that were the case, there would be an $\varepsilon \in \mathbb{R}_{++}$ with

$$\lim_{n \rightarrow \infty} \min_{p_a \in \Delta(Y)} \mathbb{E}_{p_a} [u(a, y)] + \frac{1}{\lambda_n} R(p_a || q_a^{\theta_n}) = \mathbb{E}_{q_a^\theta} [u(a, y)] - \varepsilon. \quad (7)$$

For every $n \in \mathbb{N}$, let $p_a^n \in \Delta(Y)$ be an arbitrary element of

$$\operatorname{argmin}_{p_a \in \Delta(Y)} \mathbb{E}_{p_a} [u(a, y)] + \frac{1}{\lambda_n} R(p_a || q_a^{\theta_n}).$$

Since Y is compact, by Theorem 15.11 in Aliprantis and Border (2013) so is $\Delta(Y)$, and therefore, we can assume (by restricting to a subsequence) that p_a^n converges to some $\hat{p}_a \in \Delta(Y)$. By equation (7) and the fact that $\lim_n p_a^n = \hat{p}_a$, we have

$$\mathbb{E}_{\hat{p}_a} [u(a, y)] \leq \mathbb{E}_{q_a^\theta} [u(a, y)] - \varepsilon.$$

Therefore,

$$\begin{aligned} & \int_0^{\bar{u}} 1 - \hat{p}_a \left(\left\{ y \in Y : u(a, y) \leq x + \min_{y \in Y} u(a, y) \right\} \right) dx + \frac{3}{4} \varepsilon \\ & \leq \int_0^{\bar{u}} 1 - q_a^\theta \left(\left\{ y \in Y : u(a, y) \leq x + \min_{y \in Y} u(a, y) \right\} \right) dx. \end{aligned} \quad (8)$$

Claim 1. *There exist $M \in \mathbb{R}$ and $L \in \mathbb{R}_{++}$ such that*

$$\hat{p}_a(\{y \in Y : u(a, y) \leq M - L\}) - q_a^\theta(\{y \in Y : u(a, y) \leq M\}) \geq \frac{\varepsilon}{2\bar{u}}. \quad (9)$$

Proof of the Claim. Suppose that for every $M \in \mathbb{R}$ and $L \in \mathbb{R}_{++}$ equation (9) does not hold. Then for every $L \in \mathbb{R}_{++}$

$$\begin{aligned}
& \int_0^{\bar{u}} 1 - \hat{p}_a \left(\left\{ y \in Y : u(a, y) \leq x + \min_{y \in Y} u(a, y) \right\} \right) dx \\
& \geq \int_L^{\bar{u}+L} 1 - \hat{p}_a \left(\left\{ y \in Y : u(a, y) \leq x + \min_{y \in Y} u(a, y) - L \right\} \right) dx - L \\
& \geq \int_L^{\bar{u}} 1 - q_a^\theta \left(\left\{ y \in Y : u(a, y) \leq x + \min_{y \in Y} u(a, y) \right\} \right) - \frac{\varepsilon}{2\bar{u}} dx - L \\
& \quad + \int_{\bar{u}}^{\bar{u}+L} 1 - \hat{p}_a \left(\left\{ y \in Y : u(a, y) \leq x + \min_{y \in Y} u(a, y) - L \right\} \right) dx \\
& \geq \int_L^{\bar{u}+L} 1 - q_a^\theta \left(\left\{ y \in Y : u(a, y) \leq x + \min_{y \in Y} u(a, y) \right\} \right) dx - \varepsilon/2 - L \\
& \geq \int_0^{\bar{u}} 1 - q_a^\theta \left(\left\{ y \in Y : u(a, y) \leq x + \min_{y \in Y} u(a, y) \right\} \right) dx - \varepsilon/2 - 2L.
\end{aligned}$$

Since L can be chosen to be arbitrarily small, we have

$$\begin{aligned}
& \int_0^{\bar{u}} 1 - \hat{p}_a \left(\left\{ y \in Y : u(a, y) \leq x + \min_{y \in Y} u(a, y) \right\} \right) dx \\
& \geq \int_0^{\bar{u}} 1 - q_a^\theta \left(\left\{ y \in Y : u(a, y) \leq x + \min_{y \in Y} u(a, y) \right\} \right) dx - \varepsilon/2,
\end{aligned}$$

a contradiction with equation (8). □

The claim, in turn, implies that there exists $N \in \mathbb{N}$ such that for all $n \geq N$

$$p_a^n \left(\left\{ y \in Y : u(a, y) \leq M - \frac{L}{2} \right\} \right) - q_a^{\theta_n} \left(\left\{ y \in Y : u(a, y) \leq M - \frac{L}{2} \right\} \right) \geq \frac{\varepsilon}{4\bar{u}}.$$

But then

$$\begin{aligned}
& \min_{p_a \in \Delta(Y)} \mathbb{E}_{p_a} [u(a, y)] + \frac{1}{\lambda_n} R(p_a || q_a^{\theta_n}) \\
& = \mathbb{E}_{p_a^n} [u(a, y)] + \frac{1}{\lambda_n} R(p_a^n || q_a^{\theta_n}) \\
& \geq \min_{y \in Y} u(a, y) + \left(p_a^n \left(\left\{ y \in Y : u(a, y) \leq M - \frac{L}{2} \right\} \right) \log \frac{p_a^n \left(\left\{ y \in Y : u(a, y) \leq M - \frac{L}{2} \right\} \right)}{q_a^{\theta_n} \left(\left\{ y \in Y : u(a, y) \leq M - \frac{L}{2} \right\} \right)} \right) / \lambda_n \\
& \quad + \left(p_a^n \left(\left\{ y \in Y : u(a, y) > M - \frac{L}{2} \right\} \right) \log \frac{p_a^n \left(\left\{ y \in Y : u(a, y) > M - \frac{L}{2} \right\} \right)}{q_a^{\theta_n} \left(\left\{ y \in Y : u(a, y) > M - \frac{L}{2} \right\} \right)} \right) / \lambda_n
\end{aligned}$$

where the inequality follows from Theorem 1.24 in Liese and Vajda (1987). But the last term diverges to $+\infty$ as n goes to infinity, a contradiction with

$$\min_{p_a \in \Delta(Y)} \mathbb{E}_{p_a} [u(a, y)] + \frac{1}{\lambda_n} R(p_a || q_a^{\theta_n}) \leq \max_{y \in Y} u(a, y) < \infty.$$

■

Lemma 4. 1. For every $a \in A$, the function $G : \Delta(\Theta) \times \mathbb{R}_+ \rightarrow \mathbb{R}$ defined by

$$G(\nu, \lambda) = \int_{\Theta} \min_{p_a \in \Delta(Y)} \mathbb{E}_{p_a} [u(a, y)] + \frac{1}{\lambda} R(p_a || q_a^{\theta}) d\nu(\theta) \quad \forall \nu \in \Delta(\Theta), \forall \lambda \in \mathbb{R}_{++}$$

and

$$G(\nu, 0) = \int_{\Theta} \mathbb{E}_{q_a^{\theta}} [u(a, y)] d\nu(\theta) \quad \forall \nu \in \Delta(\Theta)$$

is continuous.

2. The correspondence $BR^{(\cdot)}(\cdot) : \mathbb{R}_+ \times \Delta(\Theta) \rightrightarrows A$ where

$$BR^{(0)}(\nu) := BR^{Seu}(\nu) \quad \forall \nu \in \Delta(\Theta)$$

is upper hemicontinuous.

Proof. (1) Fix $a \in A$. For every $\theta \in \Theta$, let $F(\theta, 0) := \mathbb{E}_{q_a^{\theta}} [u(a, y)]$ and observe that for each $\lambda \in \mathbb{R}_{++}$, by Proposition 1.4.2 in Dupuis and Ellis (2011) we have

$$F(\theta, \lambda) := \min_{p_a \in \Delta(Y)} \mathbb{E}_{p_a} [u(a, y)] + \frac{1}{\lambda} R(p_a || q_a^{\theta}) = -\frac{1}{\lambda} \log \left(\int_Y \exp(-\lambda u(a, y)) dq_a^{\theta}(y) \right).$$

Since Y is compact and u is continuous, for all $\lambda \in \mathbb{R}_{++}$ and $\theta \in \Theta$, the RHS belongs to

$$\left[\min_{y \in Y} u(a, y), \max_{y \in Y} u(a, y) \right].$$

For every $\lambda \in \mathbb{R}_{++}$, $\int_Y \exp(-\lambda u(a, \cdot))$ is a continuous and bounded function that is bounded away from 0. Moreover, by Assumption (ii) and Scheffe's theorem (see Theorem in 16.12 in Billingsley, 2008), $\theta \mapsto q_a^{\theta}$ is continuous. Therefore,

$$\theta \mapsto \int_Y \exp(-\lambda u(a, y)) dq_a^{\theta}(y)$$

is continuous by definition of the weak convergence of measures, and F is continuous by Lemma 3 (at $\lambda = 0$) and Theorem 15.7.3 in Kallenberg (1973) (at $\lambda \neq 0$).

Let $(\nu_n, \lambda_n)_{n \in \mathbb{N}} \in \Delta(\Theta) \times \mathbb{R}_{++}$ be a convergent sequence with limit (ν, λ) . Suppose first that $\lambda > 0$. Then

$$\begin{aligned} & \lim_{n \rightarrow \infty} \int_{\Theta} -\frac{\log \left(\int_Y \exp(-\lambda_n u(a, y)) dq_a^\theta(y) \right)}{\lambda_n} d\nu_n(\theta) \\ &= \int_{\Theta} -\frac{\log \left(\int_Y \exp(-\lambda u(a, y)) dq_a^\theta(y) \right)}{\lambda} d\nu(\theta) \end{aligned}$$

by Theorem 15.7.3 in Kallenberg (1973) and the joint continuity of F established above. Next, suppose that $\lambda = 0$. Then

$$\lim_{n \rightarrow \infty} \int_{\Theta} -\frac{\log \left(\int_Y \exp(-\lambda u(a, y)) dq_a^\theta(y) \right)}{\lambda_n} d\nu_n(\theta) = \int_{\Theta} \int_Y u(a, y) dq_a^\theta(y) d\nu(\theta)$$

again by Theorem 15.7.3 in Kallenberg (1973) and the joint continuity of F established above. This proves (i).

(2) Follows by (1) and Theorem 17.31 in Aliprantis and Border (2013). ■

Lemma 5. 1. For every $a \in A$, if $(\theta_n, p_a^n)_{n \in \mathbb{N}} \in (\Theta \times \Delta(Y))^{\mathbb{N}}$ is such that

$$\lim_{n \rightarrow \infty} (\theta_n, p_a^n)_{n \in \mathbb{N}} = (\bar{\theta}, \bar{p}_a)$$

and $\text{supp} p_a^n \subseteq \{y \in Y : -\ln \tilde{q}_a^\theta(y) \leq K\}$ for all $n \in \mathbb{N}$ and $\theta \in \Theta$ then

$$\lim_{n \rightarrow \infty} -\int_Y \log(\tilde{q}_a^{\theta_n}(y)) dp_a^n(y) = -\int_Y \log(\tilde{q}_a^{\bar{\theta}}(y)) d\bar{p}_a(y).$$

2. For every $a \in A$, if $(\theta_n, \theta'_n, p_a^n)_{n \in \mathbb{N}} \in (\Theta \times \Theta \times \Delta(Y))^{\mathbb{N}}$ is such that

$$\lim_{n \rightarrow \infty} (\theta_n, \theta'_n, p_a^n)_{n \in \mathbb{N}} = (\bar{\theta}, \bar{\theta}', \bar{p}_a)$$

and $\text{supp} p_a^n \subseteq \{y \in Y : -\ln \tilde{q}_a^\theta(y) \leq K\}$ for all $n \in \mathbb{N}$ and $\theta \in \Theta$ then

$$\lim_{n \rightarrow \infty} -\int_Y \log\left(\frac{\tilde{q}_a^{\theta_n}(y)}{\tilde{q}_a^{\theta'_n}(y)}\right) dp_a^n(y) = -\int_Y \log\left(\frac{\tilde{q}_a^{\bar{\theta}}(y)}{\tilde{q}_a^{\bar{\theta}'}(y)}\right) d\bar{p}_a(y).$$

Proof. By Assumption 1 (i-ii) and Scheffe's theorem (see Theorem in 16.12 in Billings-

ley, 2008), the assumptions of Theorem 15.7.3 in Kallenberg (1973) are satisfied for the sequences of integrand functions and probability measures $(\log(\tilde{q}_a^{\theta_n}(\cdot)), p_a^n)_{n \in \mathbb{N}}$ and $(\log(\tilde{q}_a^{\theta_n}(\cdot)), p_a^n)_{n \in \mathbb{N}}$. ■

The following lemma shows that under every policy, almost surely the infinite sequence of observations do not contain a realization that provides an arbitrarily large evidence against a structured model.

Lemma 6. *For every $\Pi \in A^{\mathcal{H}}$ and $\theta \in \Theta$,*

$$\mathbb{P}_{\Pi} \left(\left\{ (a_i, y_i)_{i \in \mathbb{N}} \in (A \times Y)^{\mathbb{N}} : \forall t \in \mathbb{N}, -\ln \tilde{q}_{a_t}^{\theta}(y_t) \leq K \right\} \right) = 1.$$

Proof. By Assumption 1 (i) for every $a \in A$, $\{y \in Y : -\ln \tilde{q}_a^{\theta}(y) \leq K\} \in \mathcal{B}(Y)$. By the definition of Radon-Nykodim derivative and equation (6), for every $t \in \mathbb{N}$,

$$\mathbb{P}_{\Pi} \left(\left\{ (a_i, y_i)_{i \in \mathbb{N}} \in (A \times Y)^{\mathbb{N}} : -\ln \tilde{q}_{a_t}^{\theta}(y_t) > K \right\} \right) = 0.$$

Since \mathbb{P}_{Π} is a measure, it is countably subadditive and so

$$\begin{aligned} & \mathbb{P}_{\Pi} \left(\left\{ (a_i, y_i)_{i \in \mathbb{N}} \in (A \times Y)^{\mathbb{N}} : \forall t \in \mathbb{N}, -\ln \tilde{q}_{a_t}^{\theta}(y_t) \leq K \right\} \right) \\ &= 1 - \mathbb{P}_{\Pi} \left(\left\{ (a_i, y_i)_{i \in \mathbb{N}} \in (A \times Y)^{\mathbb{N}} : \exists t, -\ln \tilde{q}_{a_t}^{\theta}(y_t) > K \right\} \right) \\ &\geq 1 - \sum_{t=1}^{\infty} \mathbb{P}_{\Pi} \left(\left\{ (a_i, y_i)_{i \in \mathbb{N}} \in (A \times Y)^{\mathbb{N}} : -\ln \tilde{q}_{a_t}^{\theta}(y_t) > K \right\} \right) = 1, \end{aligned}$$

proving the statement. ■

The following lemma shows that, on every history where the empirical action process stabilizes on α^* , and the empirical outcome distribution contingent on the actions played infinitely often converges to the true distribution, the limit LRT can be rewritten as the minimum of an α^* weighted average of the relative entropy from the true DGP.

Lemma 7. *Let $\alpha^* \in \Delta(A)$ and $(a_t, y_t)_{t \in \mathbb{N}} \in (A \times Y)^{\mathbb{N}}$ be such that there exists $\theta \in \Theta$ with $-\ln \tilde{q}_{a_t}^{\theta}(y_t) \leq K$ for all $t \in \mathbb{N}$. For every $t \in \mathbb{N}$, set $h_t = (a^t, y^t)$, and let $\hat{\theta}(h_t)$ and $r(h_t)$ be two arbitrary elements of $\operatorname{argmax}_{\theta \in \Theta} \prod_{\tau=1}^t \tilde{p}_{a_{\tau}}^{\theta}(y_{\tau})$ and $\operatorname{argmax}_{p \in N(\Theta)} \prod_{\tau=1}^t \tilde{p}_{a_{\tau}}(y_{\tau})$, respectively. If*

$$\lim_{t \rightarrow \infty} \left(\alpha_t(h_t), (p_a^{h_t})_{a \in \operatorname{supp} \alpha^*} \right) = \left(\alpha^*, (p_a^*)_{a \in \operatorname{supp} \alpha^*} \right)$$

then

$$\lim_{t \rightarrow \infty} \sum_{a \in A} \sum_{\tau=1}^t \mathbb{I}_{\{a\}}(a_\tau) \int_Y \log \left(\frac{\tilde{r}_a(h_t)(y)}{\tilde{q}_a^{\hat{\theta}(h_t)}(y)} \right) dp_a^{h_t}(y) / t = \min_{\theta \in \Theta} \sum_{a \in A} \alpha^*(a) R(p_a^* || q_a^\theta).$$

Moreover,

$$\lim_{t \rightarrow \infty} \lambda^\beta(h_t) = \lim_{t \rightarrow \infty} \frac{t}{\beta_t} \cdot \min_{\theta \in \Theta} \sum_{a \in A} \alpha^*(a) R(p_a^* || q_a^\theta).$$

Proof. By assumption of the lemma, we have

$$\begin{aligned} & \sum_{a \in A} \sum_{\tau=1}^t \mathbb{I}_{\{a\}}(a_\tau) \int_Y \log \left(\frac{\tilde{r}_a(h_t)(y)}{\tilde{q}_a^{\hat{\theta}(h_t)}(y)} \right) dp_a^{h_t}(y) / t \\ &= \frac{\sum_{a \in A} \sum_{\tau=1}^t \mathbb{I}_{\{a\}}(a_\tau) \left(\int_Y \log(\tilde{r}_a(h_t)(y)) dp_a^{h_t}(y) - \int_Y \log(\tilde{q}_a^{\hat{\theta}(h_t)}(y)) dp_a^{h_t}(y) \right)}{t}. \end{aligned}$$

Let $a \in \text{supp} \alpha^*$. By Assumption 2 (i), we have

$$\int_Y -\log(\tilde{r}_a(h_t)(y)) dp_a^{h_t}(y) \leq 0 \quad \forall t \in \mathbb{N}.$$

Also, take any subsequence h_{t_n} in which $r_a(h_t)$ converges to some r_a

$$0 \leq \int_Y -\log(\tilde{r}_a(y)) dp_a^*(y) \leq \liminf_{n \rightarrow \infty} \int_Y -\log(\tilde{r}_a(h_{t_n})(y)) dp_a^{h_{t_n}}(y)$$

where the first inequality follows from Gibbs inequality and the second since by Assumption 2 (ii), there exists $K' \in \mathbb{R}_{++}$ such that $-\log(\tilde{r}_a(h_{t_n})(y)) \geq -K'$, p_a^* -almost surely and so we can apply Lemma 3.2 in Serfozo (1982). Therefore, we have

$$\lim_{t \rightarrow \infty} \sum_{a \in A} \sum_{\tau=1}^t \mathbb{I}_{\{a\}}(a_\tau) \int_Y -\log(\tilde{r}_a(h_t)(y)) dp_a^{h_t}(y) / t = 0.$$

So

$$\begin{aligned}
& \lim_{t \rightarrow \infty} \sum_{a \in A} \sum_{\tau=1}^t \mathbb{I}_{\{a\}}(a_\tau) \int_Y \log \left(\frac{\tilde{r}_a(h_t)(y)}{\tilde{q}_a^{\hat{\theta}(h_t)}(y)} \right) dp_a^{h_t}(y) / t \\
&= - \lim_{t \rightarrow \infty} \sum_{a \in A} \sum_{\tau=1}^t \mathbb{I}_{\{a\}}(a_\tau) \int_Y \log \left(\tilde{q}_a^{\hat{\theta}(h_t)}(y) \right) dp_a^{h_t}(y) / t \\
&= - \lim_{t \rightarrow \infty} \min_{\theta \in \Theta} \sum_{a \in A} \sum_{\tau=1}^t \mathbb{I}_{\{a\}}(a_\tau) \int_Y \log \left(\tilde{q}_a^\theta(y) \right) dp_a^{h_t}(y) / t.
\end{aligned}$$

Therefore the first part of the result follows from Lemma 5 and Theorem 17.31 in Aliprantis and Border (2013).

For the last part of the claim observe that

$$\begin{aligned}
\lim_{t \rightarrow \infty} \lambda^\beta(h_t) &= \lim_{t \rightarrow \infty} \sum_{a \in A} \sum_{\tau=1}^t \mathbb{I}_{\{a\}}(a_\tau) \int_Y \log \left(\frac{\tilde{r}_a(h_t)(y)}{\tilde{q}_a^{\hat{\theta}(h_t)}(y)} \right) dp_a^{h_t}(y) / \beta_t \\
&= \lim_{t \rightarrow \infty} \frac{t}{\beta_t} \cdot \sum_{a \in A} \sum_{\tau=1}^t \mathbb{I}_{\{a\}}(a_\tau) \int_Y \log \left(\frac{\tilde{r}_a(h_t)(y)}{\tilde{q}_a^{\hat{\theta}(h_t)}(y)} \right) dp_a^{h_t}(y) / t \\
&= \lim_{t \rightarrow \infty} \frac{t}{\beta_t} \cdot \min_{\theta \in \Theta} \sum_{a \in A} \alpha^*(a) R(p_a^* || q_a^\theta)
\end{aligned}$$

where the first equality follows from Lemma 2 and the third equality follows from the first part of the statement. \blacksquare

Define the set $\Theta^\varepsilon(a)$ as all parameters at most ε away from a relative entropy minimizer given action a ,

$$\Theta^\varepsilon(a) = \{\theta \in \Theta : \text{there exists } \theta' \in \Theta(a) \text{ with } \|\theta - \theta'\|_\infty \leq \varepsilon\}. \quad (10)$$

Lemma 8. *For every $a \in A$ and $\varepsilon \in (0, 1)$, there exists $\epsilon \in (0, 1)$ such that we have*

$$\lim_{t \rightarrow \infty} \inf_{\substack{\theta \in \Theta \setminus \Theta^\varepsilon(a) \\ \theta' \in \Theta^\epsilon(a)}} \int_Y \log \left(\frac{\tilde{q}_a^\theta(y)}{\tilde{q}_a^{\theta'}(y)} \right) dp_a^{h_t}(y) \geq \epsilon \quad \mathbb{P}_{\Pi^a} \text{ a.s.}$$

Proof. By the Glivenko-Cantelli Theorem (see, e.g., Theorem 2.4.1 and Example 2.1.3 in Van der Vaart and Wellner, 1997 for the version that applies to $Y \subseteq \mathbb{R}^k$) we have $\lim_{n \rightarrow \infty} p_a^{h_n} = p_a^*$, \mathbb{P}_{Π^a} a.s. Consider an history where this convergence realizes, and

suppose by contradiction that there is a convergent sequence $(\theta_n, \theta'_n)_{n \in \mathbb{N}}$ in $\Theta \times \Theta$ with limit (θ, θ') , $\theta_n \in \Theta^{1/n}(a)$, $\theta'_n \in \Theta \setminus \Theta^\varepsilon(a)$ and

$$\int_Y \log(\tilde{q}_a^{\theta_n}(y)) - \log(\tilde{q}_a^{\theta'_n}(y)) dp_a^{h_n}(y) \leq \frac{1}{n}.$$

But we would have

$$\theta \in \Theta(a), \quad \theta' \in \Theta \setminus \Theta^{\varepsilon/2}(a),$$

and by Lemma 5

$$\int_Y \log(\tilde{q}_a^\theta(y)) - \log(\tilde{q}_a^{\theta'}(y)) dp_a^*(y) = 0$$

a contradiction. ■

Lemma 9. *For every $a \in A$,*

$$\limsup_{k \rightarrow \infty} \min_{\theta \in \Theta} \min_{p_a \in \Delta(Y)} \int_Y u(a, y) dp_a(y) + \frac{R(p_a || q_a^\theta)}{k} = \min_{y \in \cup_{\theta \in \Theta} \text{supp} q_a^\theta} u(a, y).$$

Proof. Let $\hat{y} \in \operatorname{argmin}_{y \in \cup_{\theta \in \Theta} \text{supp} q_a^\theta} u(a, y)$. If $\max_{y \in Y} u(a, y) = u(a, \hat{y})$ the statement is trivially true, so suppose that $\max_{y \in Y} u(a, y) > u(a, \hat{y})$. By Assumption 1 (i) we have that

$$\inf_{\theta \in \Theta} q_a^\theta(B_\varepsilon(\hat{y})) > 0 \quad \forall \varepsilon \in \mathbb{R}_{++}.$$

Otherwise, by the compactness of Θ and the continuity of q_a^θ in θ guaranteed by the Scheffe's theorem (see Theorem 16.12 in Billingsley (2008)), the portmanteau theorem (see, e.g., Theorem 11.1.1 Dudley, 2018) would imply that there exists $\hat{\theta} \in \Theta$ with $q_a^{\hat{\theta}}(B_{\varepsilon/2}(\hat{y})) = 0$. But then, since there exists $\bar{\theta} \in \Theta$ with $\hat{y} \in \text{supp} q_a^{\bar{\theta}} = \text{supp} p_a^*$, and so $p_a^*(B_{\varepsilon/2}(\hat{y})) > 0$, we would obtain a contradiction with $p_a^* \sim q_a^{\hat{\theta}}$. Fix

$$\bar{\varepsilon} \in \left(0, \frac{\max_{y \in Y} u(a, y) - u(a, \hat{y})}{2}\right).$$

Since $u(a, \cdot)$ is continuous, there exists ε such that

$$y \in B_\varepsilon(\hat{y}) \implies u(a, y) \leq u(a, \hat{y}) + \bar{\varepsilon}.$$

Then, for all $\theta \in \Theta$

$$\begin{aligned}
u(a, \hat{y}) &\leq \min_{p_a \in \Delta(Y)} \int_Y u(a, y) dp_a + \frac{R(p_a || q_a^\theta)}{k} \\
&= -\frac{1}{k} \log \left(\int_Y \exp(-ku(a, y)) dq_a^\theta(y) \right) \\
&\leq -\frac{1}{k} \log \left(\frac{\exp(-k(u(a, \hat{y}) + \bar{\varepsilon})) \inf_{\hat{\theta} \in \Theta} q_a^{\hat{\theta}}(B_\varepsilon(\hat{y}))}{+ \left(1 - \inf_{\hat{\theta} \in \Theta} q_a^{\hat{\theta}}(B_\varepsilon(\hat{y}))\right) \exp(-k \max_{y \in Y} u(a, y))} \right)
\end{aligned}$$

where the equality follows from Proposition 1.4.2 in Dupuis and Ellis (2011). Moreover, the last term converges to $u(a, \hat{y}) + \bar{\varepsilon}$ as k goes to infinity by a simple application of L'Hôpital's rule. Since $\bar{\varepsilon} < \frac{\max_{y \in Y} u(a, y) - u(a, \hat{y})}{2}$ was arbitrarily chosen, and the last term does not depend on Θ this proves the desired uniformity of the convergence. Then the statement follows from Lemma 6 in Dunford and Schwartz (1988). \blacksquare

Proof of Proposition 1. Let (u, a, Y) be an exogenous decision problem and $\varepsilon \in \mathbb{R}_{++}$. We start by showing that there exists $c \in \mathbb{R}_{++}$ such that $\beta_t = ct$ is ε -safe and ε -consistent under almost correct specification. This is done by first deriving a $c \in \mathbb{R}_{++}$ such that ε -safety is satisfied, and then showing that there exists a δ that delivers ε -consistency under almost correct specification. Safety is trivially satisfied by every policy if

$$\max_{\bar{a} \in A} \min_{y \in Y} u(\bar{a}, y) = \min_{a \in A, y \in Y} u(a, y),$$

so in that case pick an arbitrary $c \in \mathbb{R}_{++}$. Suppose instead that we have

$$\max_{\bar{a} \in A} \min_{y \in Y} u(\bar{a}, y) > \min_{a \in A, y \in Y} u(a, y).$$

Let $\hat{P} \subseteq \Delta(Y)^A$ be the set of p^* that satisfy Assumption 1 jointly with Θ , and define

$$\underline{A}(p^*) := \left\{ a' \in A : \max_{\bar{a} \in A} \min_{y \in Y} u(\bar{a}, y) > \mathbb{E}_{p_a^*}[u(a', y)] + \frac{\varepsilon}{2} \right\}.$$

Claim 2. *There exists $\varphi^* > 0$ such that for every $\Pi \in A^{\mathcal{H}}$ and $p^* \in \hat{P}$,*

$$\mathbb{P}_\Pi \left(\left\{ \liminf_{t \rightarrow \infty} \frac{\sum_{i=1}^t u(a_i, y_i)}{t} - \max_{a \in A} \min_{y \in Y} u(a, y) - \varepsilon < 0 \right\} \cap \{ \limsup_{t \rightarrow \infty} \alpha_t(h_t)(a') < \varphi^*, \forall a' \in \underline{A}(p^*) \} \right) = 0.$$

That is, almost surely the payoff is at most ε -lower than the safe guarantee if the

actions whose objective expected performance is lower than the guarantee are played sufficiently rarely (i.e., each of them has an average frequency lower than φ^*).

Proof of the Claim. Consider the stochastic process defined by

$$X_t = u(\Pi(h_{t-1}), y_t) - \mathbb{E}_{p_{\Pi(h_{t-1})}^*} [u(\Pi(h_{t-1}), y)] \quad \forall t \in \mathbb{N}$$

with the sequence of sigma-algebras $(\mathcal{F}_t)_{t \in \mathbb{N}}$ generated by the stochastic process of histories $(h_t)_{t \in \mathbb{N}}$. The stochastic process is not i.i.d., as previous utility realizations affect current period choices. Nevertheless it is a martingale difference sequence, as u is continuous in y on the compact Y , so $\mathbb{E}[|X_t|] \leq 2 \max_{a,y} |u(a, y)| < \infty$ and $\mathbb{E}[X_t | \mathcal{F}_{t-1}] = 0$ by equation (6). A fortiori, $(X_t)_{t \in \mathbb{N}}$ is a mixingale difference sequence, and by the strong law of large numbers for mixingale sequences (see Theorem 2.7 in Hall and Heyde, 2014 for the version that applies here), we have

$$\lim_{n \rightarrow \infty} \frac{\sum_{t=1}^n X_t}{n} = 0 \quad \mathbb{P}_{\Pi}\text{-a.s.}$$

so that

$$\begin{aligned} \liminf_{t \rightarrow \infty} \frac{\sum_{i=1}^t u(a_i, y_i)}{t} &= \liminf_{t \rightarrow \infty} \sum_{i=1}^t \frac{X_i + \mathbb{E}_{p_{a_t}^*} [u(a_t, \cdot)]}{t} \\ &\geq \left(1 - \limsup_{t \rightarrow \infty} \alpha_t(\underline{A}(p^*))\right) \left(\max_{\bar{a} \in A} \min_{y \in Y} u(\bar{a}, y) - \frac{\varepsilon}{2}\right) \\ &\quad + \limsup_{t \rightarrow \infty} \alpha_t(\underline{A}(p^*)) \min_{a \in A, y \in Y} u(a, y) \\ &\geq \left(1 - \sum_{a \in \underline{A}(p^*)} \limsup_{t \rightarrow \infty} \alpha_t(a)\right) \max_{\bar{a} \in A} \min_{y \in Y} u(\bar{a}, y) - \frac{\varepsilon}{2} \\ &\quad + \sum_{a \in \underline{A}(p^*)} \limsup_{t \rightarrow \infty} \alpha_t(a) \min_{a \in A, y \in Y} u(a, y) \\ &\geq \left(1 - |A| \max_{a \in \underline{A}(p^*)} \limsup_{t \rightarrow \infty} \alpha_t(a)\right) \max_{\bar{a} \in A} \min_{y \in Y} u(\bar{a}, y) - \frac{\varepsilon}{2} \\ &\quad + \left(|A| \max_{a \in \underline{A}(p^*)} \limsup_{t \rightarrow \infty} \alpha_t(a)\right) \min_{a \in A, y \in Y} u(a, y) \end{aligned}$$

and therefore the claim follows from setting

$$\frac{\varepsilon}{2(\max_{\bar{a} \in A} \min_{y \in Y} u(\bar{a}, y) - \min_{a \in A, y \in Y} u(a, y)) |A|} = \varphi^*.$$

□

Claim 3. *There exists $\bar{\lambda} \in \mathbb{R}_{++}$ such that if $\lambda \geq \bar{\lambda}$ then for every exogenous $p^* \in \hat{P}$, $a' \in \underline{A}(p^*)$, $\nu \in \Delta(\Theta)$, we have $a' \notin BR^\lambda(\nu)$.*

That is, if the agent is sufficiently misspecification concerned, they do not play actions that can perform worse than the safe guarantee.

Proof of the Claim. Observe that if $\underline{A}(p^*) \neq \emptyset$, then by Assumption 1 (i) for all $\theta \in \Theta$, there is $y \in \text{supp } q_{a'}^\theta$ with $u(a', y) \leq \max_{\bar{a} \in A} \min_{y \in Y} u(\bar{a}, y) - \frac{\varepsilon}{2}$. But then the claim follows from Lemma 9. □

Claim 4. *There exists $J \in (0, 1)$ such that for every exogenous $p^* \in \hat{P}$, $a' \in \underline{A}(p^*)$, $\mu \in \Delta(\Theta)$, and $\lambda \in \mathbb{R}_+$,*

$$\mu(\{\theta : R(p^*||q^\theta) > J\}) \leq J \implies a' \notin BR^\lambda(\mu). \quad (11)$$

That is, if the beliefs are sufficiently concentrated on the parameters that are close to the true DGP, and under the true DGP a' performs worse than the safe guarantee, a' cannot be chosen regardless of the level of misspecification concern.

Proof of the Claim. Observe that given Claim 3, the statement immediately holds for $\lambda > \bar{\lambda}$. Suppose by contradiction that equation (11) does not hold true. This means that there exists a convergent $(p_n^*, \mu_n, \lambda_n)_{n \in \mathbb{N}} \in \hat{P} \times \Delta(\Theta) \times [0, \bar{\lambda}]$ and $a' \in A$ with

$$\mu_n \left(\left\{ \theta : R(p_n^*||q^\theta) > \frac{1}{n} \right\} \right) \leq \frac{1}{n}, \quad a' \in \underline{A}(p_n^*), \text{ and } a' \in BR^{\lambda_n}(\mu_n).$$

By the lower semicontinuity of R and the fact that $R(p||q) = 0$ if and only if $p = q$, (see, e.g., Lemma 1.4.3 in Dupuis and Ellis (2011)), as well as Lemma 4 this in turn implies that there exists $\theta \in \Theta$ with

$$a' \in \underline{A}(q_\theta) \text{ and } \mathbb{E}_{q_{a'}^\theta}[u(a', y)] \geq \max_{\bar{a} \in A} \min_{y \in Y} u(\bar{a}, y),$$

a contradiction. □

By Claim 4 every policy is safe if $\min_{\theta \in \Theta} R(p^*||q^\theta) < J$. Therefore, let $c = \frac{J}{2\bar{\lambda}}$, $\beta_t := ct$, and Π be a β -optimal policy. Take an arbitrary $p^* \in \hat{P}$ with $\min_{\theta \in \Theta} R(p^*||q^\theta) \geq J$. By Lemmas 6 and 7 we have that

$$\liminf_{t \rightarrow \infty} \lambda^\beta(h_t) \geq \frac{\min_{\theta \in \Theta} R(p^*||q^\theta)}{c} \geq \frac{J}{c} = 2\bar{\lambda} \quad \mathbb{P}_\Pi\text{-a.s.}$$

Then by Claim 3, we have that for all $a' \in \underline{A}(p^*)$

$$\limsup_{t \rightarrow \infty} \alpha_t(a') = 0 < \varphi^* \quad \mathbb{P}_\Pi\text{-a.s.}$$

But then the ε -safety of β follows by Claim 2.

Since Θ is compact, for every $\epsilon \in (0, 1)$ we can pick $\delta < 0$ such that for all $p^* \in \hat{P}$,

$$\min_{\theta \in \Theta} R(p^* || q^\theta) < \delta$$

implies that for $\theta \in \Theta^\epsilon(a)$

$$\mathbb{E}_{p^*}[u(a', y)] - \min_{p \in \Delta(Y)} \mathbb{E}_p[u(a', y)] - \frac{R(p || q_a^\theta)}{\lambda} \leq \frac{\varepsilon}{4} \quad \forall a' \in A, \forall \lambda \in [0, 2\delta].$$

But then $\beta = ct$ is ε -consistent with this δ by Lemmas 6, 7, and 8.

We next prove a stronger claim that the one stated in the proposition. We show that there is an exogenous decision problem (u, a, Y) such that if the concern for misspecification of the agent is such that

$$\lambda(h_t) = o\left(\frac{LRT(h_t, \Theta)}{t}\right) \quad \mathbb{P}_\Pi\text{-a.s.}$$

then the decision rule is not $\frac{1}{10}$ -safe. That, is we do not assume that the concern for misspecification satisfies the specific ratio structure of equation (2), but that is just asymptotically smaller than the average likelihood ratio test across all sample paths. Suppose that

$$A = \{1, -1, 0\} \quad \text{and} \quad Y = \{-1, 1\}.$$

The utility function is $u(a, y) = ay$. Since the decision problem is exogenous, each model p considered by the agent is described by $p_a(1)$ for some arbitrary $a \in A$. With this, let $\Theta = \{0.9, 0.4\}$, $p_a^*(1) = 0.6$, and

$$\mu(0.9) = \frac{1}{2} = \mu(0.4).$$

Let $N(\Theta) = [0, 1]$, i.e., the unstructured models include all the exogenous data-

generating processes. We have

$$\max_{\bar{a} \in A} \min_{y \in Y} u(\bar{a}, y) = \min_{y \in Y} u(0, y) = 0.$$

However, by the Strong Law of Large Numbers it follows that \mathbb{P}_Π -almost surely

$$\lim_{t \rightarrow \infty} \frac{\sum_{\tau=1}^t \mathbb{I}_{\{1\}}(y_\tau)}{t} = 0.6.$$

Therefore, by Lemma 2 we have that

$$\lim_{t \rightarrow \infty} \frac{LRT(h_t, \Theta)}{t} = R(0.6||0.4) \quad \mathbb{P}_\Pi\text{-a.s.}$$

and so

$$\lim_{t \rightarrow \infty} \lambda^\beta(h_t) = 0 \quad \mathbb{P}_\Pi\text{-a.s.}$$

Moreover, for the constant function $\phi(\varepsilon) = \frac{1}{2}$ for all $\varepsilon \in \mathbb{R}_{++}$ the prior is ϕ -positive on Θ in the sense of Fudenberg, Lanzani, and Strack (2022a), and by their Lemma 1

$$\mu(0.4|h_t) \rightarrow 1 \quad \mathbb{P}_\Pi\text{-a.s.}$$

But then by the upperhemicontinuity of $BR^{(\cdot)}(\cdot)$ established in Lemma 4

$$\liminf_{t \rightarrow \infty} \frac{\sum_{i=1}^t u(a_i, y_i)}{t} = -0.2 < 0 = \max_{\bar{a} \in A} \min_{y \in Y} u(\bar{a}, y) \quad \mathbb{P}_\Pi\text{-a.s.}$$

proving the desired result.

Finally, we prove an additional stronger claim that the one stated in the proposition. We show that there is an exogenous decision problem (u, a, Y) such that if the concern for misspecification of the agent is such that

$$o(\lambda(h_t)) = \frac{LRT(h_t, \Theta)}{t} \quad \mathbb{P}_\Pi\text{-a.s.}$$

then the decision rule is not $\frac{1}{10}$ -consistent. Again we do not assume that the concern for misspecification satisfies the specific ratio structure of equation (2), but that is just asymptotically larger than the average likelihood ratio test across all sample paths.

Let $\delta \in (0, 0.4)$ and suppose

$$A = \{1, -1, 0\} \quad \text{and} \quad Y = \{-1, 1\}.$$

The utility function is $u(a, y) = ay$. Again, since the decision problem is exogenous, each model p considered by the agent is described by $p_a(1)$ for some arbitrary $a \in A$. With this, let $\Theta = \{0.6, 0.4\}$ and $p_a^*(1) = 0.6 + \delta$. Let $N(\Theta) = [0, 1]$, i.e., the unstructured models include all the exogenous data-generating processes.

Let $\bar{\lambda}$ be such that $\{0\} = BR^\lambda(\mu)$ for all $\lambda \geq \bar{\lambda}$ and $\mu \in \Delta(\Theta)$. Such a $\bar{\lambda}$ exists because for $a \in \{-1, 1\}$ and $\theta \in \Theta$

$$\lim_{\lambda \rightarrow \infty} \min_{p_a \in \Delta(Y)} \mathbb{E}_{p_a}[u(a, y)] + \frac{1}{\lambda} R(p_a || q_a^\theta) = -1.$$

Let

$$C_t = \left\{ h_t \in \mathcal{H}_t : p^{h_t}(1) \geq 0.6, R(p^{h_t} || 0.6) \geq \frac{R(0.6 + \delta || 0.6)}{2} \right\}.$$

For every $h_t \in C_t$, by Lemmas 2, 6, and 7

$$\lim_{t \rightarrow \infty} \frac{LRT(h_t, \Theta)}{t} \geq \frac{R(0.6 + \delta || 0.6)}{2}$$

so that $\lambda^\beta(h_t)$ is diverging to $+\infty$ and it is eventually larger than $\bar{\lambda}$. But by Sanov's Theorem (see, e.g., Theorem 2.2.1 in Dupuis and Ellis, 2011) the set C_t has a probability converging to 1, so the result follows. \blacksquare

The next lemma simplifies the task of proving that an action b cannot be a β -limit. Indeed, rather than looking at the probability that b becomes the limit action with positive probability under the candidate optimal policy, it is enough to check that if the agent plays b forever, they eventually want to change action. However, this verification under the simpler policy must be performed with respect to all priors supported over Θ and “time-shifts” of the actual concern for misspecification.

Lemma 10. *Let $b \in A$. If*

$$\mathbb{P}_{\Pi^b}[b \in BR^{\frac{k+LRT(h_\tau, \Theta)}{2\beta_\tau+n}}(\nu(\cdot|h_\tau)), \forall \tau \geq 0] = 0$$

for every prior belief ν with support Θ , $k \in \mathbb{R}_+$, and $n \in \mathbb{N}_0$, then b is not a β -limit action.

Proof. Suppose by way of contradiction that b is a β -limit action. Thus, since for every policy $\Pi \in A^{\mathcal{H}}$

$$\mathbb{P}_{\Pi}[\sup\{t: a_t \neq b\} < \infty] \leq \sum_{t=0}^{\infty} \sum_{h_t \in \mathcal{H}_t} \mathbb{P}_{\Pi} \left[b \in BR^{\frac{LRT(h_t)}{2\beta_t}}(\nu(\cdot|h_t)), \forall \tau \geq t | h_t \right] \mathbb{P}_{\tilde{\Pi}}[h_t],$$

there are a β -optimal policy $\tilde{\Pi} \in A^{\mathcal{H}}$, $t \in \mathbb{N}_0$, and $(a^t, y^t) \in \mathcal{H}_t$ with $\mathbb{P}_{\tilde{\Pi}}[(a^t, y^t)] > 0$ such that with positive probability $\tilde{\Pi}$ prescribes b after (a^t, y^t) in every future period. Define $\nu = \mu_0(\cdot|(a^t, y^t))$, and notice that by Assumption 1 (i)

$$\text{supp } \nu = \text{supp } \mu_0 = \Theta.$$

As the evolution of beliefs and misspecification concern under Π^b , i.e., the policy that plays b in every period, is the same as under $\tilde{\Pi}$ for every history where the agent continues to play b , we have that

$$\begin{aligned} & \mathbb{P}_{\tilde{\Pi}}[b = \tilde{\Pi}(h_{\tau}) \text{ for all } \tau > t | (a^t, y^t)] > 0 \\ \implies & \mathbb{P}_{\Pi^b}[b \in BR^{\frac{LRT((a^t, y^t), \Theta) + LRT(h_{\tau}, \Theta)}{2\beta_{\tau+t}}}(\nu(\cdot|(a^{\tau}, y^{\tau}))) \text{ for all } \tau > 0] > 0. \end{aligned}$$

However, the latter equals zero by the assumption of the lemma, which establishes that b cannot be a β -limit action. ■

Proof of Proposition 2. Suppose that a^* is not a self-confirming equilibrium. By Lemma 10 it is enough to show that for every prior ν with support Θ , every $k \in \mathbb{R}_+$, and every $n \in \mathbb{N}_0$,

$$\inf \left\{ t \in \mathbb{N}_0 : a^* \notin BR^{\frac{k + LRT((a^t, y^t), \Theta)}{2\beta_{n+t}}}(\nu(\cdot|(a^t, y^t))) \right\} < \infty \quad \mathbb{P}_{\Pi_{a^*}}\text{-a.s.}$$

By the strong law of large numbers (see, e.g., Theorem 8.3.5 in Dudley, 2018),

$$\lim_{t \rightarrow \infty} p_{a^*}^{h_t} = p_{a^*}^* \quad \mathbb{P}_{\Pi_{a^*}}\text{-a.s.}$$

Therefore, by the assumptions of the proposition and Gibbs' inequality

$$\lim_{t \rightarrow \infty} \prod_{\tau=1}^t p_{a^*}^{h_t}(y_{\tau}) = \lim_{t \rightarrow \infty} \max_{\theta \in \Theta} \prod_{\tau=1}^t q_{a^*}^{\theta}(y_{\tau}) = \lim_{t \rightarrow \infty} \max_{p \in N(\Theta)} \prod_{\tau=1}^t p_{a^*}(y_{\tau}) \quad \mathbb{P}_{\Pi_{a^*}}\text{-a.s.}$$

So, $\lim_{t \rightarrow \infty} LRT(h_t, \Theta) = 0$, $\mathbb{P}_{\Pi_{a^*}}$ almost surely and also

$$\lim_{t \rightarrow \infty} \frac{k + LRT(h_t, \Theta)}{2\beta_{n+t}} = 0 \quad \mathbb{P}_{\Pi_{a^*}}\text{-a.s.}$$

With this, as by Lemma 8 for every $\varepsilon \in \mathbb{R}_{++}$ we have

$$\nu(B_\varepsilon(\{\theta \in \Theta : p_{a^*}^\theta = p_{a^*}^*\}) | h_t) \rightarrow 1, \quad \mathbb{P}_{\Pi_{a^*}}\text{-a.s.}$$

the desired conclusion follows from Lemma 4. ■

Proof of Theorem 1. We start with the preliminary observation that by Lemma 6, $-\ln \tilde{q}_{a_t}^\theta(y_t) \leq K$ for all $t \in \mathbb{N}$ and $\theta \in \Theta$, $\mathbb{P}_{\Pi_{a^*}}$ -almost surely. This will allow us to invoke Lemma 7 in all the various cases.

1) Suppose that a^* is not a Berk-Nash equilibrium. By Lemma 10, it is enough to show that for every prior ν with support Θ , every $k \in \mathbb{R}_+$, and every $n \in \mathbb{N}_0$, if the agent plays a^* in every period, then almost surely they arrive to a history h_t such that

$$a^* \notin BR^{\frac{k+LRT(h_t, \Theta)}{2\beta_{n+t}}}(\nu(\cdot | h_t)).$$

By the Glivenko-Cantelli Theorem (see, e.g., Theorem 2.4.1 and Example 2.1.3 in Van der Vaart and Wellner, 1997 for the version that applies to $Y \subseteq \mathbb{R}^k$), $p_{a^*}^{h_t} \rightarrow p_{a^*}^*$, $\mathbb{P}_{\Pi_{a^*}}$ -a.s.

Then, by Lemma 7, since $t = o(\beta_t)$, we have

$$\lim_{t \rightarrow \infty} \frac{k + LRT(h_t, \Theta)}{2\beta_{n+t}} = \lim_{t \rightarrow \infty} \frac{k}{2\beta_{n+t}} + \lim_{t \rightarrow \infty} \frac{\beta_t}{\beta_{n+t}} \lambda^\beta(h_t) = 0 \quad \mathbb{P}_{\Pi_{a^*}}\text{-a.s.}$$

By Assumption 1 (ii), the assumptions of Berk (1966), page 54, are satisfied, and we have that for every $\varepsilon \in \mathbb{R}_{++}$,

$$\nu(\Theta^\varepsilon(a^*) | h_t) \rightarrow 1, \quad \mathbb{P}_{\Pi_{a^*}}\text{-a.s.}$$

Therefore, since Θ is compact, $\left(\frac{k+LRT(h_t, \Theta)}{2\beta_{n+t}}, \nu(\cdot | h_t)\right)_{t \in \mathbb{N}}$ admits $\mathbb{P}_{\Pi_{a^*}}$ almost surely a subsequence convergent to $(0, \nu^*)$ for some $\nu^* \in \Delta(\Theta(a^*))$. With this, the result follows from Lemma 4.

2) Suppose that

$$a^* \notin BR^{Meu} \left(\left\{ p \in \Delta(Y)^A : \exists \theta \in \Theta, \forall a \in A, q_a^\theta \gg p_a \right\} \right).$$

By Lemma 10, it is enough to show that for every prior ν with support Θ , every $k \in \mathbb{R}_+$, and every $n \in \mathbb{N}_0$, if the agent plays a^* in every period, then almost surely they arrive to a history h_t such that $a^* \notin BR^{\frac{k+LRT(h_t, \Theta)}{2\beta_{n+t}}}(\nu(\cdot|h_t))$. By the Glivenko-Cantelli Theorem (see, e.g., Theorem 2.4.1 and Example 2.1.3 in Van der Vaart and Wellner, 1997 for the version that applies to $Y \subseteq \mathbb{R}^k$),

$$p_{a^*}^{h_t} \rightarrow p_{a^*}^* \quad \mathbb{P}_{\Pi_{a^*}}\text{-a.s.}$$

Then, by Lemma 7, since $\beta_t = o(t)$ and $\min_{\theta \in \Theta} R(p_{a^*}^* || q_{a^*}^\theta) > 0$, we have

$$\lim_t \lambda^\beta(h_t) = \infty \quad \mathbb{P}_{\Pi_{a^*}}\text{-a.s.}$$

Moreover, we have

$$\begin{aligned} \lim_{t \rightarrow \infty} \frac{k + LRT(h_t, \Theta)}{2\beta_{n+t}} &= \lim_{t \rightarrow \infty} \frac{k}{2\beta_{n+t}} + \lim_{t \rightarrow \infty} \frac{LRT(h_t, \Theta)}{LRT(h_{t+n}, \Theta)} \lambda^\beta(h_{t+n}) \\ &\geq \lim_{t \rightarrow \infty} \frac{t}{t+n} \lambda^\beta(h_{t+n}) = \infty \quad \mathbb{P}_{\Pi_{a^*}}\text{-a.s.} \end{aligned}$$

where the inequality follows from the first part of Lemma 7. By Assumption 1 (i) for all $\theta, \theta' \in \Theta$ and $a \in A$ we have

$$q_a^\theta \sim q_a^{\theta'}.$$

So we have

$$\left\{ p \in \Delta(Y)^A : \exists \theta \in \Theta, \forall a \in A, q_a^\theta \gg p_a \right\} = \left\{ p \in \Delta(Y)^A : \forall \theta \in \Theta, \forall a \in A, q_a^\theta \gg p_a \right\}.$$

Therefore, by Lemma 9 for all $a \in A$ we have that $\mathbb{P}_{\Pi_{a^*}}$ almost surely

$$\limsup_{t \rightarrow \infty} \min_{\theta \in \Theta} \min_{p_a \in \Delta(Y)} \int_Y u(a, y) dp_a + \frac{2\beta_{n+t} R(p_a || q_a^\theta)}{k + LRT(h_t, \Theta)} = \min_{y \in \cup_{\theta \in \Theta} \text{supp} q_a^\theta} u(a, y).$$

But since by Assumption 1 (i) for all $t \in \mathbb{N}$, $\mu(\cdot|h_t) \subseteq \Theta$, $\mathbb{P}_{\Pi_{a^*}}$ almost surely, we

have

$$\lim_{t \rightarrow \infty} \mathbb{E}_{\mu(\cdot|h_t)} \left[\min_{p_a \in \Delta(Y)} \int_Y u(a, y) dp_a + \frac{2\beta_{n+t} R(p_a || q_a^\theta)}{k + LRT(h_t, \Theta)} \right] = \min_{y \in \cup_{\theta \in \Theta} \text{supp} q_a^\theta} u(a, y) \quad \mathbb{P}_{\Pi_{a^*}}\text{-a.s.}$$

With this, the result follows from the finiteness of the action space.

3) Let $\hat{\theta} \in \text{argmin}_{\theta \in \Theta} R(p_{a^*}^* || q_{a^*}^\theta)$. Suppose that for every $\nu \in \Delta(\Theta)$ with

$$\text{supp} \nu \subseteq \text{argmin}_{\theta \in \Theta} R(p_{a^*}^* || q_{a^*}^\theta),$$

we have $a^* \notin BR^{R(p_{a^*}^* || q_{a^*}^{\hat{\theta}})/c}(\nu)$. By Lemma 10, it is enough to show that for every prior ν with support Θ , every $k \in \mathbb{R}_+$, and every $n \in \mathbb{N}_0$, if the agent plays a^* in every period, then almost surely they arrive to a history h_t such that

$$a^* \notin BR^{\frac{k + LRT(h_t, \Theta)}{2\beta_{n+t}}}(\nu(\cdot|h_t)).$$

By the Glivenko-Cantelli Theorem (see, e.g., Theorem 2.4.1 and Example 2.1.3 in Van der Vaart and Wellner, 1997 for the version that applies to $Y \subseteq \mathbb{R}^k$),

$$p_{a^*}^{h_t} \rightarrow p_{a^*}^* \quad \mathbb{P}_{\Pi_{a^*}}\text{-a.s.}$$

By Lemma 7

$$\lim_{t \rightarrow \infty} \frac{k + LRT(h_t, \Theta)}{2\beta_{n+t}} = \lim_{t \rightarrow \infty} \frac{k}{2\beta_{n+t}} + \lim_{t \rightarrow \infty} \frac{\beta_t}{\beta_{n+t}} \lambda^\theta(h_t) = R(p_{a^*}^* || q_{a^*}^{\hat{\theta}}) / c \quad \mathbb{P}_{\Pi_{a^*}}\text{-a.s.}$$

By Assumption 1 (ii), the assumptions of Berk (1966), page 54, are satisfied, and we have that for every $\varepsilon \in \mathbb{R}_{++}$, $\nu(\Theta_\varepsilon(a^*)|h_t) \rightarrow 1$, $\mathbb{P}_{\Pi_{a^*}}$ -almost surely. As Θ is compact, by Theorem 15.11 in Aliprantis and Border (2013) so is $\Delta(\Theta)$. This means that a subsequence of $(\nu(\cdot|h_t))_{t \in \mathbb{N}}$ converges to a $\bar{\nu} \in \Delta(\Theta)$ with

$$\text{supp} \bar{\nu} \subseteq \Theta(a^*).$$

Therefore the result follows from the upperhemicontinuity of $BR^{(\cdot)}(\cdot)$ established in Lemma 4. ■

Lemma 11. *For every $c \in \mathbb{R}_{++}$ the function $\alpha \mapsto \min_{\theta \in \Theta} \sum_{a \in A} \alpha(a) R(p_a^* || q_a^\theta) / c$ is continuous and the correspondence $\Theta(\cdot) : \Delta(A) \rightarrow 2^\Theta$ is upper hemicontinuous.*

Proof. We first show that the function

$$(\alpha, \theta) \mapsto \sum_{a \in A} \alpha(a) R(p_a^* || q_a^\theta) / c \quad (12)$$

is continuous. Fix an $a \in A$ and let $(\theta_n)_{n \in \mathbb{N}}$ be a sequence that converges to θ . By Assumption 1 (ii), $\tilde{q}_a^{\theta_n}$ is converging pointwise to \tilde{q}_a^θ . Then

$$|R(p_a^* || q_a^{\theta_n}) - R(p_a^* || q_a^\theta)| = \left| \int_Y \log \left(\frac{\tilde{q}_a^\theta(y)}{\tilde{q}_a^{\theta_n}(y)} \right) dp_a^*(y) \right|$$

and observe that the integrand on the right-hand side is dominated by a constant by Assumption 1 (i). Therefore, by the dominated convergence theorem $|R(p_a^* || q_a^{\theta_n}) - R(p_a^* || q_a^\theta)|$ converges to 0. Since A is finite and the function in equation (12) is linear in α , we have obtained the desired continuity. With this, the statement follows from Theorem 17.31 in Aliprantis and Border (2013).

Proof of Proposition 3. Consider the following three-player game. The action sets are $A_1 = \Delta(A)$, $A_2 = \Delta(\Theta)$, $A_3 = \mathbb{R}_+$ with arbitrary elements denoted as α, ν, λ . The utility functions are

$$\begin{aligned} U_1(\alpha, \nu, \lambda) &= \begin{cases} \sum_{a \in A} \alpha(a) \int_{\Theta} \min_{p_a \in \Delta(Y)} \mathbb{E}_{p_a} [u(a, y)] + \frac{1}{\lambda} R(p_a || q_a^\theta) d\nu(\theta) & \lambda \neq 0 \\ \sum_{a \in A} \alpha(a) \int_{\Theta} \mathbb{E}_{q_a^\theta} [u(a, y)] d\nu(\theta) & \lambda = 0, \end{cases} \\ U_2(\alpha, \nu, \lambda) &= - \int_{\Theta} \sum_{a \in A} \alpha(a) R(p_a^* || q_a^\theta) d\nu(\theta), \\ U_3(\alpha, \nu, \lambda) &= - \left(\lambda - \min_{\bar{\theta} \in \Theta} \sum_{a \in A} \alpha(a) R(p_a^* || q_a^{\bar{\theta}}) / c \right)^2. \end{aligned}$$

Observe that for the purpose of finding the equilibria of this game, it is without loss of generality to limit the actions of player 3 to $[0, \bar{\lambda}]$ where

$$\bar{\lambda} = \max_{\alpha \in \Delta(A)} \min_{\bar{\theta} \in \Theta} \sum_{a \in A} \alpha(a) R(p_a^* || q_a^{\bar{\theta}}) / c < \infty,$$

where the inequality holds by Assumption 1 (i). Therefore, since the compactness of Θ implies that also $\Delta(\Theta)$ is compact by Theorem 15.11 in Aliprantis and Border (2013) all the three action sets are compact. Moreover, they are clearly convex.

The utility function U_1 is jointly continuous in its second and third argument by

Lemma 4. Moreover, U_2 is trivially continuous in its first and third argument and U_3 is continuous in its first and second argument by Lemma 11. Therefore the game is better-reply secure (see Reny, 1999, page 1033). Moreover, U_1 and U_2 are respectively linear in A_1 and A_2 while U_3 is concave in A_3 .

Therefore, by Theorem 3.1 and Footnote 8 in Reny (1999) this game admits a pure-strategy equilibrium $(\alpha^*, \nu^*, \lambda^*)$. But observe that

$$\begin{aligned} \lambda^* &\in \operatorname{argmax}_{\lambda \in \mathbb{R}_+} - \left(\lambda - \min_{\bar{\theta} \in \Theta} \sum_{a \in A} \alpha^*(a) R(p_a^* || q_a^{\bar{\theta}}) / c \right)^2 \\ \implies \lambda^* &= \min_{\bar{\theta} \in \Theta} \sum_{a \in A} \alpha^*(a) R(p_a^* || q_a^{\bar{\theta}}) / c, \end{aligned}$$

$$\alpha^* \in \operatorname{argmax}_{\alpha \in \Delta(A)} U_1(\alpha, \nu^*, \lambda^*) \implies \alpha^* \in \Delta(BR^{\lambda^*}(\nu^*)),$$

and

$$\begin{aligned} \nu^* &\in \operatorname{argmax}_{\nu \in \Delta(\Theta)} - \int_{\Theta} \sum_{a \in A} \alpha^*(a) R(p_a^* || q_a^{\theta}) d\nu(\theta) \\ \implies \operatorname{supp} \nu^* &\subseteq \operatorname{argmin}_{\theta \in \Theta} \sum_{a \in A} \alpha^*(a) R(p_a^* || q_a^{\theta}) \\ \implies \nu^* &\in \Delta(\Theta(\alpha^*)). \end{aligned}$$

Therefore, α^* is a mixed c -robust equilibrium sustained by the belief ν^* and the concern for misspecification λ^* . ■

Proof of Theorem 2. We start with the preliminary observation that by Lemma 6, $-\ln \tilde{q}_{a_t}^{\theta}(y_t) \leq K$ for all $t \in \mathbb{N}$ and $\theta \in \Theta$, \mathbb{P}_{Π} -almost surely. This will allow us to invoke Lemma 7. Observe that $(\alpha_t)_{t \in \mathbb{N}}$ satisfies the following differential inclusion: for all $a \in A$, $t \in \mathbb{N}$, $h_t \in \mathcal{H}_t$, and $h_{t+1} \in \mathcal{H}_{t+1}$ such that $h_{t+1} \succ h_t$

$$\alpha_{t+1}(h_{t+1})(a) \in \left\{ \alpha_t(h_t)(a) + \frac{1}{t+1} (\mathbb{I}_{\{a'\}}(a) - \alpha_t(h_t)(a)) : a' \in BR^{\lambda^{\theta}(h_t)}(\mu(\cdot | h_t)) \right\}.$$

Set $\tau_0 = 0$ and $\tau_t = \sum_{i=1}^t \frac{1}{i}$ for all $t \in \mathbb{N}$. The continuous-time interpolation of α_t is the function $w : \mathbb{R}_+ \rightarrow \Delta(A)$

$$w(\tau_t + l) = \begin{cases} \alpha_t + l \frac{\alpha_{t+1} - \alpha_t}{\tau_{t+1} - \tau_t}, & \forall t \in \mathbb{N}, \forall l \in [0, \frac{1}{t+1}] \\ \alpha_1 & t = 0, \forall l \in [0, 1]. \end{cases} \quad (13)$$

For every $\alpha \in \Delta(A)$, let

$$\chi_\alpha = \left\{ \alpha' \in \Delta \left(BR^{\min_{\theta \in \Theta} \sum_{a \in A} \alpha(a) R(p_a^* || q_a^\theta) / c} (\Delta(\Theta(\alpha))) \right) \right\} \subseteq \Delta(A).$$

We use the theory of stochastic approximation for differential inclusions (Benaim, Hofbauer, and Sorin, 2005 and Esponda, Pouzo, and Yamamoto, 2021a) to show that (13) can be approximated by a solution to

$$\dot{\alpha}_t \in \chi_{\alpha_t} - \alpha_t. \quad (14)$$

A solution over $[0, T]$, $T \in \mathbb{R}_{++}$, to the differential inclusion (14) with initial point $\hat{\alpha} \in \Delta(A)$ is a mapping $\alpha_{(\cdot)} : [0, T] \rightarrow \Delta(A)$ that is absolutely continuous over compact intervals such that $\alpha_0 = \hat{\alpha}$ and (14) is satisfied for almost every t . Let $S_{\hat{\alpha}}^T$ be the set of the solutions to (14) over $[0, T]$, $T \in \mathbb{R}_{++}$, with initial conditions $\hat{\alpha} \in \Delta(A)$. Since by Lemmas 11 and 4 $\alpha \mapsto \min_{\theta \in \Theta} \sum_{a \in A} \alpha(a) R(p_a^* || q_a^\theta) / c$ is continuous and $\Theta(\cdot)$, $BR^{(\cdot)}(\cdot)$ are upper hemicontinuous,

$$\alpha \mapsto \Delta \left(BR^{\min_{\theta \in \Theta} \sum_{a \in A} \alpha(a) R(p_a^* || q_a^\theta) / c} (\Delta(\Theta(\alpha))) \right) \quad (15)$$

is upper hemicontinuous. To see this, we show that it has a closed graph. Since $\Delta(A)$ is compact, this is enough by, e.g., Proposition E.3 in Ok (2011). Let

$$(\alpha_n, \alpha'_n)_{n \in \mathbb{N}} \in (\Delta(A) \times \Delta(A))^{\mathbb{N}}$$

be such that

$$\alpha_n \in \Delta \left(BR^{\min_{\theta \in \Theta} \sum_{a \in A} \alpha'_n(a) R(p_a^* || q_a^\theta) / c} (\Delta(\Theta(\alpha'_n))) \right) \quad \forall n \in \mathbb{N}$$

and $\lim_{n \rightarrow \infty} (\alpha_n, \alpha'_n) = (\alpha, \alpha')$. By finiteness of A , we can take (possibly truncating some initial elements of the sequence) $(\alpha_n, \alpha'_n)_{n \in \mathbb{N}}$ to be such that $\text{supp} \alpha \subseteq \text{supp} \alpha_n$ for all $n \in \mathbb{N}$. Then for every $\bar{a} \in \text{supp} \alpha$ there is $(\nu_n^{\bar{a}})_{n \in \mathbb{N}}$ such that $\nu_n^{\bar{a}} \in \Delta(\Theta(\alpha'_n))$ and

$$\bar{a} \in BR^{\min_{\theta \in \Theta} \sum_{a \in A} \alpha'_n(a) R(p_a^* || q_a^\theta) / c} (\nu_n^{\bar{a}}) \quad \forall n \in \mathbb{N}.$$

Since Θ is compact so is $\Delta(\Theta)$ by Theorem 15.11 in Aliprantis and Border (2013), and by restricting to a subsequence we can take $\nu_n^{\bar{a}}$ to be convergent to some $\nu^{\bar{a}} \in \Delta(\Theta)$.

Since Θ is upper hemicontinuous, $\nu^{\bar{a}} \in \Theta(\alpha')$. Since $BR^{(\cdot)}(\cdot)$ is upper hemicontinuous,

$$\bar{a} \in BR^{\min_{\theta \in \Theta} \sum_{a \in A} \alpha'(a) R(p_a^* || q_a^\theta)/c}(\nu^{\bar{a}}) \subseteq BR^{\min_{\theta \in \Theta} \sum_{a \in A} \alpha'(a) R(p_a^* || q_a^\theta)/c}(\Theta(\alpha'))$$

showing that (α, α') belongs to the graph of correspondence (15). Therefore, as χ_α is also convex- and closed-valued, a solution to (14) exists by Theorem 2.1.4 in Aubin and Cellina (2012), i.e., S_α^T is nonempty for every $T \in \mathbb{R}_{++}$ and $\hat{\alpha} \in \Delta(A)$. Let $S^T = \cup_{\hat{\alpha} \in \Delta(A)} S_\alpha^T$.

We next establish that the continuous-time interpolation of $(\alpha_t(h_t))_{t \in \mathbb{N}}$ defined in (13) can in the long run be approximated arbitrarily well by a solution to (14). Observe that w is Lipschitz continuous of order 1 as for all history sequences $(h_t)_{t \in \mathbb{N}} \in \times_{t \in \mathbb{N}} \mathcal{H}_t$ with $h_{t+1} \succ h_t$ for all $t \in \mathbb{N}$,

$$\frac{\|\alpha_{t+1}(h_{t+1}) - \alpha_t(h_t)\|_\infty}{\tau_{t+1} - \tau_t} \leq \frac{1/(t+1)}{1/(t+1)} = 1 \quad \forall t \in \mathbb{N}. \quad (16)$$

Therefore w is absolutely continuous (see, e.g., Proposition 7 in Royden and Fitzpatrick, 1988), and α_t is uniformly bounded because it takes values in $\Delta(A)$. Let $\Upsilon = \{\alpha - \alpha' : \alpha, \alpha' \in \Delta(A)\}$ and for all $\varepsilon \in \mathbb{R}_+$ and $\alpha' \in \Delta(A)$,

$$M_\varepsilon(\alpha') = \left\{ \nu \in \Delta(\Theta) : \int_\Theta \sum_{a \in A} \alpha'(a) R(p_a^* || q_a^\theta) d\nu(\theta) \leq \varepsilon + \min_{\theta \in \Theta} \sum_{a \in A} \alpha'(a) R(p_a^* || q_a^\theta) \right\}.$$

By Esponda, Pouzo, and Yamamoto, 2021a, Part 1a of the proof of Theorem 2 (observe that Assumption 1 (i-ii) implies their Assumption 2 (ii-iii)), $M_{(\cdot)}(\cdot)$ is upper hemicontinuous. We define

$$F : \mathbb{R}_+ \times \mathbb{R}_+ \times \Delta(A) \rightrightarrows \Upsilon$$

by

$$F(\varepsilon, \varepsilon', \alpha) = \left\{ \iota \in \left[\Delta \left(\cup_{\lambda' \in B_{\varepsilon'}} (\min_{\theta \in \Theta} \sum_{a \in A} \alpha(a) R(p_a^* || q_a^\theta)/c) \cap [0, \frac{2K}{c}] BR^{\lambda'}(\Delta(M_\varepsilon(\alpha))) \right) - \alpha \right] \right\}.$$

Observe that $F(0, 0, \alpha) = \chi_\alpha - \alpha$. Moreover, we now show that F has a closed graph, so it is upper hemicontinuous. Let

$$(\iota_n, \varepsilon_n, \varepsilon'_n, \alpha_n)_{n \in \mathbb{N}} \in (\Upsilon \times \mathbb{R}_+ \times \mathbb{R}_+ \times \Delta(A))^{\mathbb{N}}$$

be such that $\iota_n \in F(\varepsilon_n, \varepsilon'_n, \alpha_n)$ for all $n \in \mathbb{N}$ and

$$\lim_{n \rightarrow \infty} (\iota_n, \varepsilon_n, \varepsilon'_n, \alpha_n) = (\iota, \varepsilon, \varepsilon', \alpha).$$

Since A is finite, it is without loss of generality to take $\iota_n(a) > -\alpha_n(a)$ for all $n \in \mathbb{N}$ and for all a for which $\iota(a) > -\alpha(a)$. Then for all \hat{a} such that $\iota(\hat{a}) > -\alpha(\hat{a})$, there is a sequence $(\nu_n^{\hat{a}}, \lambda_n^{\hat{a}})_{n \in \mathbb{N}} \in (\Delta(\Theta) \times [0, 2K/c])^{\mathbb{N}}$ such that

$$\nu_n^{\hat{a}} \in M_{\varepsilon_n}(\alpha'_n), \lambda_n^{\hat{a}} \in B_{\varepsilon'_n} \left(\min_{\theta \in \Theta} \sum_{a \in A} \alpha_n(a) R(p_a^* || q_a^\theta) / c \right), \text{ and } \hat{a} \in BR^{\lambda_n^{\hat{a}}}(\nu_n^{\hat{a}}).$$

Since $\Delta(\Theta)$ and $[0, 2K/c]$ are compact by restricting to a subsequence we can take $(\nu_n^{\hat{a}}, \lambda_n^{\hat{a}})_{n \in \mathbb{N}}$ to be convergent to some $(\nu^{\hat{a}}, \lambda^{\hat{a}}) \in \Delta(\Theta) \times [0, 2K/c]$. Since $M_{(\cdot)}(\cdot)$ is upper hemicontinuous $\nu^{\hat{a}} \in M_\varepsilon(\alpha)$. Since

$$\hat{\alpha} \mapsto \min_{\theta \in \Theta} \sum_{a \in A} \hat{\alpha}(a) R(p_a^* || q_a^\theta) / c$$

is continuous by Lemma 11,

$$\lambda \in B_{\varepsilon'} \left(\min_{\theta \in \Theta} \sum_{a \in A} \alpha(a) R(p_a^* || q_a^\theta) / c \right).$$

Since $BR^{(\cdot)}(\cdot)$ is upper hemicontinuous,

$$\hat{a} \in BR^{\lambda^{\hat{a}}}(\nu^{\hat{a}}) \subseteq \left\{ BR^{\hat{\lambda}}(\hat{\nu}) : \hat{\nu} \in M_\varepsilon(\alpha), \hat{\lambda} \in B_{\varepsilon'} \left(\min_{\theta \in \Theta} \sum_{a \in A} \alpha(a) R(p_a^* || q_a^\theta) / c \right) \cap \left[0, \frac{2K}{c} \right] \right\}$$

showing that $(\iota, \varepsilon, \varepsilon', \alpha)$ belongs to the graph of the correspondence (15).

With this, $F(0, 0, \cdot) + (\cdot) : \Delta(A) \rightarrow \Delta(A)$ satisfies Hypothesis 1.1 in Benaim, Hofbauer, and Sorin (2005), as it is clearly compact- and convex-valued. Moreover, by Theorem 1 in Esponda, Pouzo, and Yamamoto (2021a) and Lemma 7 we have that \mathbb{P}_Π -almost surely, if $\lim_{t \rightarrow \infty} \alpha_t = \alpha^*$, we eventually have $\mu(\cdot | h_t) \in M_\varepsilon(\alpha^*)$ and

$$\lambda^\beta(h_t) \in B_{\varepsilon'} \left(\min_{\theta \in \Theta} \sum_{a \in A} \alpha^*(a) R(p_a^* || q_a^\theta) / c \right)$$

for all $(\varepsilon, \varepsilon') \in \mathbb{R}_{++}^2$. Thus, there is a sequence $(\hat{\varepsilon}_t)_{t \in \mathbb{N}} \in \mathbb{R}_{++}^{\mathbb{N}}$ converging to 0 with

$$\chi_{\alpha_t} - \alpha_t \in F(\hat{\varepsilon}_t, \hat{\varepsilon}_t, \alpha_t).$$

Fix $T \in \mathbb{N}$ and define the flow operator $G : C(\mathbb{R}, \Delta(A)) \times \mathbb{R} \rightarrow C(\mathbb{R}, \Delta(A))$ as

$$G^t(f)(s) = f(s+t) \quad \forall f \in C(\mathbb{R}, \Delta(A)), \forall s \in \mathbb{R}, \forall t \in \mathbb{R}.$$

We now show that every limit point of $(G^t(w))_{t \in \mathbb{N}}$ is in S^T . This argument borrows extensively from the proofs Theorem 4.2 in Benaim, Hofbauer, and Sorin (2005) and Theorem 2 in Esponda, Pouzo, and Yamamoto (2021a). However they cannot be directly applied, because the interpolated process w we consider is not a perturbed solution in the sense of Benaim, Hofbauer, and Sorin (2005). Indeed, it may not be possible to find an α that *jointly* justifies a_t as a best reply to beliefs in $\Theta(\alpha)$ and the concern for misspecification $\min_{\theta \in \Theta} \sum_{a \in A} \alpha(a) R(p_a^* || q_a^\theta) / c$, as perturbation of the empirical frequency α_{t-1} in different directions may be needed for the concern and the belief. Nevertheless, the core of their arguments can be adapted leveraging the upper-hemicontinuity of F established above.

Since w is uniformly continuous by equation (16), the family $(G^t(w))_{t \in \mathbb{N}}$ is equicontinuous, and thus it is relatively compact in the topology of uniform convergence over compact sets by the Arzela-Ascoli theorem (see Willard, 2012 Theorem 43.15 for the version with a noncompact domain). The topology of uniform convergence over compact sets is metrizable since $\Delta(A)$ is metrizable and \mathbb{R} is open (see Theorem 1.14b in Simon, 2020), and so there exists a limit point $z = \lim_{t_n} G^{t_n}(w)$. Define

$$m(t) = \max \{k \in \mathbb{N} : \tau_k \leq t\}$$

and for all $s \in \mathbb{R}$, $v(s) = w'(s) = \alpha_{m(s)+1} - \alpha_{m(s)} \in F(\hat{\varepsilon}_{m(s)}, \hat{\varepsilon}_{m(s)}, \alpha_{m(s)})$, and $v_n(s) = v(t_n + s)$ so

$$\begin{aligned} z(T) - z(0) &= \lim_{t_n} (G^{t_n}(w)(T) - G^{t_n}(w)(0)) \\ &= \lim_{t_n} (w(T + t_n) - w(t_n)) = \lim_{n \rightarrow \infty} \int_0^T v_n(s) ds. \end{aligned}$$

Since $(v_n)_{n \in \mathbb{N}}$ is uniformly bounded, it is bounded in $L^2([0, T], \mathbb{R}^A, Leb)$. By the Banach-Alaoglu theorem (see Theorem 6.21 in Aliprantis and Border, 2013), (by restricting to a subsequence) we can take $(v_n)_{n \in \mathbb{N}}$ to be a weakly-convergent subsequence with limit $v^* \in L^2([0, T], \mathbb{R}^A, Leb)$. By Mazur's lemma (see Corollary V.3.14 in Dunford and Schwartz (1988)), there exist a function $N : \mathbb{N} \rightarrow \mathbb{N}$ and a sequence of positive

weights $(\rho_n(n), \dots, \rho_{N(n)}(n))_{n \in \mathbb{N}}$ with $\sum_{i=n}^{N(n)} \rho_i(n) = 1$ for all $n \in \mathbb{N}$ such that if we define

$$\bar{v}_n = \sum_{i=n}^{N(n)} \rho_i(n) v_i,$$

then \bar{v}_n converges with respect to the $L^2([0, T], \mathbb{R}^A, Leb)$ norm, and thus almost surely, to v^* .

Let $\tau \in [0, T]$ be such that the convex combination of the elements of $(v_n)_{n \in \mathbb{N}}$ is converging to v^* at τ . For every $t \in [0, T]$ and $n \in \mathbb{N}$, define

$$\gamma_n(t) = \hat{\varepsilon}_{m(t_n+t)} + \|w(t_n+t) - \alpha_{m(t_n+t)}\|$$

and

$$w_n(t) = w(t_n+t).$$

Observe that by definition of w , $(\hat{\varepsilon}_t)_{t \in \mathbb{N}}$, and z ,

$$\lim_{n \rightarrow \infty} \gamma_n(t) = 0 \text{ and } \lim_{n \rightarrow \infty} w_n(t) = z(t).$$

But then, by the upperhemicontinuity of F , for every $\varepsilon \in \mathbb{R}_{++}$ there exists N_ε such that for $n \geq N_\varepsilon$, $F(\gamma_n(t), \gamma_n(t), w_n(t)) \subseteq B_\varepsilon(F(0, 0, z(t)))$, where the latter set is closed and convex. But since $v_n(t) \in F(\gamma_n(t), \gamma_n(t), w_n(t))$, also

$$\bar{v}_n(t) \in F(\gamma_n(t), \gamma_n(t), w_n(t)) \subseteq B_\varepsilon(F(0, 0, z(t))).$$

Therefore, $v^* \in (F(0, 0, z(t)))$. Since the fact that v_n is weakly convergent to v^* implies by definition that

$$\lim_{n \rightarrow \infty} \int_0^T v_n(s) ds = \lim_{n \rightarrow \infty} \int_0^T v^*(s) ds$$

we have that $z \in S^T$.

Therefore, by (ii) \implies (i) of Theorem 4.1 in Benaim, Hofbauer, and Sorin (2005) (see Esponda, Pouzo, and Yamamoto, 2021b for the slightly corrected version used here)

$$\lim_{t \rightarrow \infty} \inf_{\tilde{\alpha} \in S^T} \sup_{0 \leq s \leq T} \|w(t+s) - \tilde{\alpha}_s\| = 0 \quad \mathbb{P}_\Pi\text{-a.s. for all } T \in \mathbb{N}. \quad (17)$$

With this, we can replicate an argument from Fudenberg, Lanzani, and Strack (2022b) to rule out convergence to non equilibria. If $\alpha^* \in \Delta(A)$ is not a mixed c -robust

equilibrium, there is $a \in A$ with $\alpha^*(a) > 0$ and $\delta_a \notin \chi_{\alpha^*}$. Since $\chi(\cdot) = F(0, 0, \cdot) + (\cdot)$ and F has a closed graph and maps into the compact Υ , there exists $D \in \mathbb{R}_{++}$ such that for all $\alpha' \in B_D(\alpha^*)$, $\alpha'(a) - \max_{\hat{\alpha} \in \chi_{\alpha'}} \hat{\alpha}(a) > \alpha^*(a)/2$. Therefore, for every initial condition $\bar{\alpha} \in B_D(\alpha^*)$ and every solution of (14), $\alpha(a)$ decreases at rate at least $\alpha^*(a)/4$ until it leaves $B_D(\alpha^*)$. So for every initial condition $\bar{\alpha} \in B_D(\alpha^*)$ and every solution, the differential inclusion leaves $B_D(\alpha^*)$ before time

$$T^* := \frac{D + \alpha^*(a)}{\alpha^*(a)/4}.$$

With this, we can prove that $(\alpha_t)_{t \in \mathbb{N}}$ does not converge to α^* on a sample path on which the convergence of equation (17) happens. Since the set of such sample paths has probability 1 under policy Π , this fact concludes the proof. Suppose by contradiction that on one of such paths α_t converges to α^* . Therefore, we can choose $\hat{T} \in \mathbb{N}$ such that on that sample path $\alpha_t \in B_{D/2}(\alpha^*)$ for all $t \geq \hat{T}$ and

$$\inf_{\tilde{\alpha} \in S^{T^*}} \sup_{0 \leq s \leq T^*} \|w(\hat{T} + s) - \tilde{\alpha}_s\| \leq D/4. \quad (18)$$

Take any $\tilde{\alpha} \in S^{T^*}$ with $\sup_{0 \leq s \leq T^*} \|w(\hat{T} + s) - \tilde{\alpha}_s\| \leq D/2$. Since $w(\hat{T}) \in B_{D/2}(\alpha^*)$, $\tilde{\alpha} \in S_{\tilde{\alpha}}^{T^*}$ for some initial condition $\bar{\alpha} \in B_D(\alpha^*)$. But then by definition of T^* the differential inclusion leaves $B_D(\alpha^*)$ by time $T^* + \hat{T}$, and by (18), α_t does not stay in $B_{D/2}(\alpha^*)$, a contradiction. \blacksquare

Proof of Corollary 1. We first show that for a sufficiently low c there is no c -robust equilibrium. Observe that by Assumption 3 (i) and Proposition 8 in Battigalli, Cerreia-Vioglio, Maccheroni, Marinacci, and Sargent (2022) for every $\alpha \in \Delta(A)$, we have

$$\Theta(\alpha) = (\theta_0^*, \theta_{1\pi}^*, \theta_{1a}^*, \theta_2^*, \theta_3^*). \quad (19)$$

Moreover, since θ^* perfectly predicts the consequences under policy 0, we have

$$\min_{\theta \in \Theta} R(p_0^* || q_0^\theta) = 0.$$

By Assumption 3 (i), and Lemma 3 in Battigalli, Cerreia-Vioglio, Maccheroni, Marinacci, and Sargent (2022), $BR^{Seu}(\Delta(\Theta(0))) = \{1\}$, and therefore 0 is not a c -robust equilibrium for any $c \in \mathbb{R}_{++}$. Since f_1 is strictly concave on \mathbb{R}_{++} , we have $\min_{\theta \in \Theta} R(p_1^* || q_1^\theta) > 0$. By Assumption 3 (ii) and Lemma 9 there exists a sufficiently

small \bar{c} such that for all $c \leq \bar{c}$,

$$BR^{\frac{\min_{\theta \in \Theta} R(p_1^* || q_1^\theta)}{c}}(\delta_{\theta^*}) = \{0\}$$

proving that there is no c -robust equilibrium if $c \leq \bar{c}$. That a mixed c -robust equilibrium exists follows from Proposition 3.³⁴

In particular, the maximal (resp. the minimal) equilibrium is defined as the α such that $\sum_{a \in A} \alpha(a) R(p_a^* || q_a^{\theta^*})$ is equal to the maximal (resp. minimal) misspecification concern λ such that $1 \in BR^\lambda(\delta_{\theta^*})$ (resp. $0 \in BR^\lambda(\delta_{\theta^*})$). Since a larger $\theta_{1\pi}^* + \theta_{1a}^*$ makes action 0 more favorable, the comparative statics follows. \blacksquare

A.2 Representation

A.2.1 Preliminaries

Let $B_0(\Sigma)$ denote the set of all real-valued Σ -measurable simple functions, and $B(\Sigma)$ the supnorm closure of $B_0(\Sigma)$. The subset of functions in $B_0(\Sigma)$ (resp. $B(\Sigma)$) that take values in $C \subseteq \mathbb{R}$ is denoted as $B_0(\Sigma, C)$ (resp. $B(\Sigma, C)$). A functional $I : \Phi \rightarrow \mathbb{R}$ defined on a nonempty subset Φ of $B(\Sigma)$ is a *niveloid* if for every $\varphi, \psi \in \Phi$

$$I(\varphi) - I(\psi) \leq \sup(\varphi - \psi).$$

It is *translation invariant* if $I(\alpha\varphi + (1 - \alpha)k\mathbb{I}_S) = I(\alpha\varphi) + (1 - \alpha)k$ for all $\alpha \in [0, 1]$, $\varphi \in \Phi$, and $k \in \mathbb{R}$ such that $\alpha\varphi + (1 - \alpha)k\mathbb{I}_S$ and $\alpha\varphi$ are in Φ . A niveloid is *normalized* if $I(k\mathbb{I}_S) = k$ for all $k \in \mathbb{R}$ such that $k\mathbb{I}_S \in \Phi$. A function $c : \Delta(S) \rightarrow \mathbb{R}_+$ is *grounded* if $c^{-1}(0) \neq \emptyset$. An event is *strongly nonnull* if for every $x, x' \in X$ with $x \succ x'$, we have $x \succ x'Ex$.

Definition 12. A tuple (u, Q, μ, ϕ) is an *average second order utility representation* of the preference relation \succsim if $u : X \rightarrow \mathbb{R}$ is a nonconstant affine function, $Q \subseteq \Delta(S)$ is a finite nonempty set, $\mu \in \Delta(Q)$, $\phi : u(X) \rightarrow \mathbb{R}$ is a strictly increasing continuous

³⁴To formally invoke Proposition 3, that requires absolute continuity with respect to the true data generating process for all $\theta \in \Theta$, restrict the parameter space to $\{\theta^*\}$. Given equation (19) every c -robust ergodic equilibrium with the reduced parameter space remains a c -robust ergodic equilibrium with the original Θ .

function, and

$$f \succsim g \iff \int_Q \phi^{-1} \left(\int_S \phi(u(f)) dq \right) d\mu(q) \geq \int_Q \phi^{-1} \left(\int_S \phi(u(g)) dq \right) d\mu(q).$$

A.2.2 Results

Our first lemma shows that the average robust control representation falls in the variational class.

Lemma 12. *Suppose that there exist a nonconstant affine function $u : X \rightarrow \mathbb{R}$, a nonempty and finite $Q \subseteq \Delta(S)$, $\mu \in \Delta(Q)$, and $(\lambda_q)_{q \in Q} \in \mathbb{R}_+^Q$ such that for all $f, g \in \mathcal{F}$*

$$f \succsim g \iff \mathbb{E}_\mu \left[\min_{p \in \Delta(S)} \mathbb{E}_p[u(f)] + \frac{R(p||q)}{\lambda_q} \right] \geq \mathbb{E}_\mu \left[\min_{p \in \Delta(S)} \mathbb{E}_p[u(g)] + \frac{R(p||q)}{\lambda_q} \right]. \quad (20)$$

Then \succsim satisfies Weak Order, Weak Certainty Independence, Continuity, Monotonicity, Uncertainty Aversion, Nondegeneracy, Weak Monotone Continuity, and admits the representation

$$f \succsim g \iff \min_{p \in \Delta(S)} \int_S \hat{u}(f) dp + \hat{c}(p) \geq \min_{p \in \Delta(S)} \int_S \hat{u}(g) dp + \hat{c}(p) \quad (21)$$

for some nonconstant affine $\hat{u} : X \rightarrow \mathbb{R}$ and a grounded, convex, and lower semicontinuous function $\hat{c} : \Delta(S) \rightarrow [0, \infty]$. Moreover, we can choose $\hat{u} = u$ and \hat{c} is such that $\hat{c}^{-1}(0) = \mathbb{E}_\mu[q]$.

Proof. We first observe that without loss of generality we can take u to be such that $0 \in \text{int } u(X)$ in the representation of equation (20). Indeed, since u is nonconstant and affine, there exists $x \in X$ with $u(x) \in \text{int } u(X)$. Define $u'(y) = u(y) - u(x)$ for all $y \in X$. Then, we have

$$\begin{aligned} f &\succsim g \\ \iff &\mathbb{E}_\mu \left[\min_{p \in \Delta(S)} \mathbb{E}_p[u(f)] + \frac{R(p||q)}{\lambda_q} \right] \geq \mathbb{E}_\mu \left[\min_{p \in \Delta(S)} \mathbb{E}_p[u(g)] + \frac{R(p||q)}{\lambda_q} \right] \\ \iff &\mathbb{E}_\mu \left[\min_{p \in \Delta(S)} \mathbb{E}_p[u'(f)] + \frac{R(p||q)}{\lambda_q} \right] \geq \mathbb{E}_\mu \left[\min_{p \in \Delta(S)} \mathbb{E}_p[u'(g)] + \frac{R(p||q)}{\lambda_q} \right] \end{aligned}$$

and $0 \in \text{int } u'(X)$.

Fix $q \in Q$. The functional $I_q : B_0(\Sigma, u(X)) \rightarrow \mathbb{R}$ defined as

$$I_q(\varphi) := \min_{p \in \Delta(S)} \int_S \varphi(s) dp + \frac{1}{\lambda_q} R(p||q) \quad \forall \varphi \in B_0(\Sigma, u(X))$$

is easily seen to be monotone, translation invariant, and concave by Theorem 11.13 in Aliprantis and Border (2013) and the concavity of the minimum. Since Q is finite,

$$\hat{I}(\varphi) = \int_Q I_q(\varphi) d\mu(q) \quad \forall \varphi \in B_0(\Sigma, u(X))$$

is well-defined and \hat{I} is monotone and concave. Let $\varphi \in B_0(\Sigma, u(X))$, $k \in u(X)$, and $\gamma \in (0, 1)$. Since u is affine and $0 \in \text{int } u(X)$, we have $\gamma\varphi + (1 - \gamma)k \in B_0(\Sigma, u(X))$, $\gamma\varphi \in B_0(\Sigma, u(X))$, and

$$\begin{aligned} \hat{I}(\gamma\varphi + (1 - \gamma)k) &= \int_Q I_q(\gamma\varphi + (1 - \gamma)k) d\mu(q) = \int_Q I_q(\gamma\varphi) + (1 - \gamma)k d\mu(q) \\ &= \int_Q I_q(\gamma\varphi) d\mu(q) + (1 - \gamma)k = \hat{I}(\gamma\varphi) + (1 - \gamma)k. \end{aligned}$$

But then, notice that

$$\int_Q \left(\min_{p \in \Delta(S)} \int_S u(f) dp + \frac{1}{\lambda_q} R(p||q) \right) d\mu(q) = \int_Q I_q(u(f)) d\mu(q) = \hat{I}(u(f))$$

where \hat{I} is monotone and translation invariant. Therefore, by Lemma 25 in Maccheroni, Marinacci, and Rustichini (2006a), $\hat{I} : B_0(\Sigma, u(X)) \rightarrow \mathbb{R}$ is a concave niveloid, and it is clearly normalized. With this, by Lemma 28 and Footnote 15 in Maccheroni, Marinacci, and Rustichini (2006a) \succsim satisfies Weak Order, Weak Certainty Independence, Continuity, Monotonicity, Uncertainty Aversion, and Nondegeneracy. Moreover, by Proposition 12 and (iii) \implies (i) of Theorem 50 in Cerreia-Vioglio, Maccheroni, Marinacci, and Montrucchio (2011), each I_q is monotone continuous, and so is \hat{I} .

We now prove it satisfies Weak Monotone Continuity. Fix $f, g \in \mathcal{F}$, $x \in X$, and $(E_i)_{i \in \mathbb{N}} \in \Sigma^{\mathbb{N}}$ with $E_1 \supseteq E_2 \supseteq \dots$, $\cap_{n \geq 1} E_n = \emptyset$, and $f \succ g$. Then $\varphi = u(f)$ and $\varphi' = u(g)$ are in $B_0(\Sigma, u(X))$, $k = u(x) \in \mathbb{R}$, and $\hat{I}(\varphi) = \hat{I}(u(f)) > \hat{I}(u(g)) = \hat{I}(\varphi')$. So by the (i) \implies (ii) of Theorem 50 in Cerreia-Vioglio, Maccheroni, Marinacci, and Montrucchio (2011) there exists $i \in \mathbb{N}$ such that $\hat{I}(k\mathbb{I}_{E_i} + \varphi\mathbb{I}_{S \setminus E_i}) > \hat{I}(\varphi')$. But observe

that

$$\hat{I}(k\mathbb{I}_{E_i} + \varphi\mathbb{I}_{S \setminus E_i}) = \hat{I}(u(x)\mathbb{I}_{E_i} + u(f)\mathbb{I}_{S \setminus E_i}) = \hat{I}(u(x_{E_i}f))$$

and $\hat{I}(\varphi') = \hat{I}(u(g))$, thus $\hat{I}(u(x_{E_i}f)) > \hat{I}(u(g))$ and $x_{E_i}f \succ g$, proving that Weak Monotone Continuity holds.

By Theorems 3 and 13 in Maccheroni, Marinacci, and Rustichini (2006a) it admits the representation in equation (21).

By the first part of the lemma we have

$$u(x) \geq u(x') \iff x \succsim x' \iff \hat{u}(x) \geq \hat{u}(x')$$

and therefore by the uniqueness up to a positive affine transformation of \hat{u} guaranteed by Corollary 5 in Maccheroni, Marinacci, and Rustichini (2006a) and the fact that every two affine functions that represent \succsim on X are positive affine transformations of each other (see, e.g., Theorem 5.11 in Kreps, 1988), we can choose $u = \hat{u}$. Finally, by (ii) \implies (iii) of Lemma 32 in Maccheroni, Marinacci, and Rustichini (2006a) for every $q \in Q$, and $k \in u(X)$, $\partial I_q(k) = \{q\}$. Let $\bar{k} \in \text{int } u(X) \neq \emptyset$ and observe that since Q is finite,

$$\begin{aligned} \lim_{\alpha \downarrow 0} \frac{\hat{I}(\bar{k} + \alpha\varphi) - \hat{I}(\bar{k})}{\alpha} &= \lim_{\alpha \downarrow 0} \frac{\mathbb{E}_\mu [I_q(\bar{k} + \alpha\varphi)] - \mathbb{E}_\mu [I_q(\bar{k})]}{\alpha} \\ &= \lim_{\alpha \downarrow 0} \mathbb{E}_\mu \left[\frac{I_q(\bar{k} + \alpha\varphi) - I_q(\bar{k})}{\alpha} \right] \\ &= \mathbb{E}_\mu \left[\lim_{\alpha \downarrow 0} \frac{I_q(\bar{k} + \alpha\varphi) - I_q(\bar{k})}{\alpha} \right] = \mathbb{E}_\mu \left[\int_S \varphi dq \right]. \end{aligned}$$

Now, applying (iii) \implies (ii) of Lemma 32 in Maccheroni, Marinacci, and Rustichini (2006a), we obtain that the unique \hat{c} identified by the choice of \hat{u} has

$$\hat{c}^{-1}(0) = \{\mathbb{E}_\mu[q]\}.$$

■

Lemma 13. *If $E \in \Sigma_{st}$ is nonnull and \succsim satisfies Weak Order, Weak Certainty Independence, Continuity, Monotonicity, Uncertainty Aversion, and Weak Monotone Continuity, then \succsim_E satisfies Weak Order, Weak Certainty Independence, Continuity, Monotonicity, Uncertainty Aversion, Nondegeneracy, and Weak Monotone Continuity.*

Proof. Let $f, g, h \in \mathcal{F}$. By Completeness of \succsim at least one between

$$fEh \succsim gEh \text{ and } gEh \succsim fEh$$

holds. Therefore, by definition of \succsim_E at least one between $f \succsim_E g$ and $g \succsim_E f$ holds.

Let $f, f', f'' \in \mathcal{F}$, with $f \succsim_E f'$ and $f' \succsim_E f''$. By definition of \succsim_E , there exist $h', h'' \in \mathcal{F}$ such that

$$fEh' \succsim f'Eh' \text{ and } f'Eh'' \succsim f''Eh''.$$

Since $E \in \Sigma_{st}$, we have

$$fEh'' \succsim f'Eh''.$$

By Transitivity of \succsim , $fEh'' \succsim f''Eh''$, and so by definition of \succsim_E , $f \succsim_E f''$.

Let $f, g \in \mathcal{F}$, $x, x' \in X$, and $\gamma \in (0, 1)$, be such that

$$\gamma f + (1 - \gamma)x \succsim_E \gamma g + (1 - \gamma)x.$$

Since $E \in \Sigma_{st}$, we have

$$(\gamma f + (1 - \gamma)x)Ex \succsim (\gamma g + (1 - \gamma)x)Ex.$$

By Weak Certainty Independence of \succsim we get

$$(\gamma f + (1 - \gamma)x')E(\gamma x + (1 - \gamma)x') \succsim (\gamma g + (1 - \gamma)x')E(\gamma x + (1 - \gamma)x').$$

But then by definition of \succsim_E , we have $\gamma f + (1 - \gamma)x' \succsim_E \gamma g + (1 - \gamma)x'$, proving that \succsim_E satisfies Weak Certainty Independence.

Let $f, g, h, h' \in \mathcal{F}$. Since $E \in \Sigma_{st}$, we have that

$$\begin{aligned} \{\gamma \in [0, 1] : \gamma f + (1 - \gamma)g \succsim_E h\} &= \{\gamma \in [0, 1] : (\gamma f + (1 - \gamma)g)Eh' \succsim hEh'\} \\ &= \{\gamma \in [0, 1] : \gamma fEh' + (1 - \gamma)gEh' \succsim hEh'\} \end{aligned}$$

and

$$\begin{aligned} \{\gamma \in [0, 1] : h \succsim_E \gamma f + (1 - \gamma)g\} &= \{\gamma \in [0, 1] : hEh' \succsim (\gamma f + (1 - \gamma)g)Eh'\} \\ &= \{\gamma \in [0, 1] : hEh' \succsim \gamma fEh' + (1 - \gamma)gEh'\} \end{aligned}$$

where the sets on the bottom lines are closed by Continuity of \succsim , proving that \succsim_E satisfies Continuity.

Let $f, g, h \in \mathcal{F}$ and $f(s) \succsim_E g(s)$ for all $s \in S$. Then, $fEh \succsim gEh$ by Monotonicity of \succsim . Therefore, by definition of \succsim_E , $f \succsim_E g$ and so \succsim_E satisfies Monotonicity.

Let $f, g, h \in \mathcal{F}$, $\gamma \in (0, 1)$ and $f \sim_E g$. Since $E \in \Sigma_{st}$, $fEh \sim gEh$ and by Uncertainty Aversion, $(\gamma f + (1 - \gamma)g)Eh = \gamma fEh + (1 - \gamma)gEh \succsim fEh$. By definition of \succsim_E , this implies that $\gamma f + (1 - \gamma)g \succsim_E f$ and so \succsim_E satisfies Uncertainty Aversion.

Since E is nonnull, there exist $f, g, h \in \mathcal{F}$ such that $fEh \succ gEh$. But then, since $E \in \Sigma_{st}$, there is no $h' \in \mathcal{F}$ with $gEh' \succ fEh'$. Therefore, by definition of \succsim_E , $f \succ_E g$ and \succsim_E satisfies Nondegeneracy.

Let $f, g, h \in \mathcal{F}$, $x \in X$, $(E_i)_{i \in \mathbb{N}} \in \Sigma^{\mathbb{N}}$ with $E_1 \supseteq E_2 \supseteq \dots$ and $\cap_{n \geq 1} E_n = \emptyset$, and $f \succ_E g$. Since $E \in \Sigma_{st}$, $fEh \succ gEh$. Moreover, $(E'_i)_{i \in \mathbb{N}}$ where $E'_i = E_i \cap E$ is such that $E'_1 \supseteq E'_2 \supseteq \dots$ and $\cap_{n \geq 1} E'_n \subseteq \cap_{n \geq 1} E_n = \emptyset$. Then $(xE'_i f)Eh = xE'_i(fEh)$ for all $i \in \mathbb{N}$ and by Weak Monotone Continuity and the fact that $fEh \succ gEh$ there exists $n_0 \in \mathbb{N}$ such that $(xE'_{n_0} f)Eh \succ gEh$. But notice that

$$(xE_{n_0} f)Eh = (xE'_{n_0} f)Eh \succ gEh$$

and therefore $xE_{n_0} f \succ_E g$, as $E \in \Sigma_{st}$. ■

Lemma 14. *Let $\Omega \times \{\rho\} \in \Sigma_{st}$ be nonnull and contain at least three disjoint nonnull events, and suppose \succsim satisfies Weak Order, Weak Certainty Independence, Continuity, Monotonicity, Uncertainty Aversion, Nondegeneracy, Weak Monotone Continuity, and the Intramodel Sure-Thing Principle. For every $f, g \in \mathcal{F}$, we have*

$$\begin{aligned} f \succsim_{\rho} g & \quad (22) \\ \iff \\ \min_{q \in \Delta(S)} \int u_{\rho}(f(s)) dq(s) + \frac{1}{\lambda_{\rho}} R(q||p_{\rho}) & \geq \min_{q \in \Delta(S)} \int u_{\rho}(g(s)) dq(s) + \frac{1}{\lambda_{\rho}} R(q||p_{\rho}) \end{aligned}$$

where u_{ρ} is a nonconstant affine function, $\lambda_{\rho} \in [0, \infty)$, and $p_{\rho} \in \Delta(S)$. Moreover, if $\Omega \times \{\rho\} \in \Sigma_{st}$ is strongly nonnull, u_{ρ} can be chosen to be the same for all such ρ and $\text{supp } p_{\rho} \subseteq \Omega \times \{\rho\}$.

Proof. By Lemma 13 \succsim_{ρ} satisfies Weak Order, Weak Certainty Independence, Continuity, Monotonicity, Uncertainty Aversion, Nondegeneracy, and Weak Monotone Con-

tinuity. We now show that for every $f, g, h, \bar{h} \in \mathcal{F}$ and $E \in \Sigma$, we have

$$fEh \succsim_{\rho} gEh \implies fE\bar{h} \succsim_{\rho} gE\bar{h}.$$

Observe that by definition of \succsim_{ρ} , $fEh \succsim_{\rho} gEh$ implies that there exists $\hat{h} \in \mathcal{F}$ such that

$$(fEh)\rho\hat{h} \succ (gEh)\rho\hat{h}.$$

But then, there exists $h' \in \mathcal{F}$ such that

$$\begin{aligned} & (fEh)\rho\hat{h} \succ (gEh)\rho\hat{h} \\ \implies & (f\{(\omega, \rho) : (\omega, \rho) \in E\}h)\rho\hat{h} \succ (g\{(\omega, \rho) : (\omega, \rho) \in E\}h)\rho\hat{h} \\ \implies & (f\{(\omega, \rho') : \rho' \in \Delta(\Omega), (\omega, \rho) \in E\}h)\rho\hat{h} \succ (g\{(\omega, \rho') : \rho' \in \Delta(\Omega), (\omega, \rho) \in E\}h)\rho\hat{h} \\ \implies & (f\{(\omega, \rho') : \rho' \in \Delta(\Omega), (\omega, \rho) \in E\}h) \succ_{\rho} (g\{(\omega, \rho') : \rho' \in \Delta(\Omega), (\omega, \rho) \in E\}h) \\ \implies & (f\{(\omega, \rho') : \rho' \in \Delta(\Omega), (\omega, \rho) \in E\}\bar{h}) \succ_{\rho} (g\{(\omega, \rho') : \rho' \in \Delta(\Omega), (\omega, \rho) \in E\}\bar{h}) \\ \implies & (f\{(\omega, \rho') : \rho' \in \Delta(\Omega), (\omega, \rho) \in E\}\bar{h})\rho h' \succ (g\{(\omega, \rho') : \rho' \in \Delta(\Omega), (\omega, \rho) \in E\}\bar{h})\rho h' \\ \implies & (f\{(\omega, \rho) : (\omega, \rho) \in E\}\bar{h})\rho h' \succ (g\{(\omega, \rho) : (\omega, \rho) \in E\}\bar{h})\rho h' \\ \implies & (fE\bar{h})\rho h' \succ (gE\bar{h})\rho h' \implies fE\bar{h} \succsim_{\rho} gE\bar{h} \end{aligned}$$

where the third, fifth, and eighth implications follow from the definition of \succsim_{ρ} , the fourth implication follows from the Intramodel Sure-Thing Principle, and the other implications only rewrite the acts involved.

Next, observe that if $E \subseteq \Omega \times \{\rho\}$ is nonnull, then there exists $f, g, h \in \mathcal{F}$ with $fEh \succ gEh$. By Structured Savage P2, this implies that $(fEh)\rho h' \succ (gEh)\rho h'$ for all $h' \in \mathcal{F}$. But this in turn means that $(fEh) \succ_{\rho} (gEh)$, so that E is nonnull for the preference \succsim_{ρ} . With this, the first part follows from Theorem 1 in Strzalecki (2011). For the second part, notice that by Theorem 3 and Lemma 30 in Maccheroni, Marinacci, and Rustichini (2006a), \succsim admits a variational representation:

$$f \succsim g \iff \min_{p \in \Delta(S)} \left(\int u(f) dp + c(p) \right) \geq \min_{p \in \Delta(S)} \left(\int u(g) dp + c(p) \right) \quad (23)$$

for some nonconstant affine $u : X \rightarrow \mathbb{R}$ and a lower semicontinuous and grounded function $c : \Delta(S) \rightarrow [0, \infty]$.

Next, assume $\Omega \times \{\rho\}$ is strongly nonnull. Notice that \succsim and \succsim_{ρ} coincide on X . Indeed, let $x \succ x'$. Since $\Omega \times \{\rho\}$ is strongly nonnull $x \succ x'\rho x$ and given that

$\Omega \times \{\rho\} \in \Sigma_{st}$ it follows that $x \succ_\rho x'$. Conversely, let $x \succsim x'$, then by equation (23) $u(x) \geq u(x')$. Since c is grounded, there exists $q^* \in \Delta(S)$ with $c(q^*) = 0$. But then

$$\begin{aligned} u(x) &\geq u(x') q^*(\Omega \times \{\rho\}) + (1 - q^*(\Omega \times \{\rho\})) u(x) \\ &\geq \min_{q \in \Delta(S)} (u(x') q(\Omega \times \{\rho\}) + (1 - q(\Omega \times \{\rho\})) u(x) + c(q)) \end{aligned}$$

that is, $x(\Omega \times \{\rho\})x \succsim x'(\Omega \times \{\rho\})x$, and $x \succsim_\rho x'$. Therefore, by the uniqueness up to a positive affine transformation of u guaranteed by Corollary 5 in Maccheroni, Marinacci, and Rustichini (2006a) and the fact that every two affine functions that represent \succsim on X are positive affine transformations of each other (see, e.g., Theorem 5.11 in Kreps, 1988), we can choose $u = u_\rho$. Suppose by way of contradiction that there exists $E \in \Sigma$ such that $E \cap (\Omega \times \{\rho\}) = \emptyset$ and $p_\rho(E) > 0$. Let $x, y \in X$ with $x \succ y$. Then,

$$u(x) > u(y) p_\rho(E) + u(x) (1 - p_\rho(E)) \geq \min_{q \in \Delta(S)} \int u(yEx) dq + \frac{1}{\lambda_\rho} R(q||p_\rho)$$

and so by equation (22), $x \succ_\rho yEx$. But since $x = x(\Omega \times \{\rho\})x$, $x = (yEx)(\Omega \times \{\rho\})x$ and $\Omega \times \{\rho\} \in \Sigma_{st}$ this would imply $x \succ x$, a contradiction to the Weak Order of \succsim . ■

Lemma 15. (i) If \succsim admits a structured average robust control representation then it admits an average second order utility representation with $\phi(\cdot) = -\exp(-\lambda(\cdot))$. (ii) If \succsim admits an average second order utility representation with $\phi(\cdot) = -\exp(-\lambda(\cdot))$, then it admits an average robust control representation.

Proof. (i) Let (u, Q, μ, λ) be a structured average robust control representation of the preference relation \succsim . By Proposition 1.4.2 in Dupuis and Ellis (2011), for all $f \in \mathcal{F}$ and $q \in \Delta(S)$

$$\min_{p \in \Delta(S)} \left(\int_S u(f) dp + \frac{1}{\lambda} R(p||q) \right) = -\frac{1}{\lambda} \log \left(\int_S \exp(-\lambda u(f)) dq \right).$$

Therefore,

$$\int_Q \min_{p \in \Delta(S)} \int_S u(f) dp + \frac{1}{\lambda} R(p||q) d\mu(q) = \int_Q -\frac{1}{\lambda} \log \left(\int_S \exp(-\lambda u(f)) dq \right) d\mu(q)$$

and the result follows from letting $\phi(\cdot) = -\exp(-\lambda(\cdot))$.

(ii) Let $(u, Q, \mu, -\exp(-\lambda(\cdot)))$ be an average second order utility representation of the preference relation \succsim . By Proposition 1.4.2 in Dupuis and Ellis (2011), for all $f \in \mathcal{F}$ and $q \in \Delta(S)$

$$-\frac{1}{\lambda} \log \left(\int_S \exp(-\lambda u(f)) \, dq \right) = \min_{p \in \Delta(S)} \left(\int_S u(f) \, dp + \frac{1}{\lambda} R(p||q) \right).$$

Therefore,

$$\int_Q \min_{p \in \Delta(S)} \int_S u(f) \, dp + \frac{1}{\lambda} R(p||q) \, d\mu(q) = \int_Q -\frac{1}{\lambda} \log \left(\int_S \exp(-\lambda u(f)) \, dq \right) \, d\mu(q).$$

■

Lemma 16. *Suppose that the assumptions of Theorem 3 hold. Let \succsim be such that for all $f, g \in \mathcal{F}$*

$$f \succsim g \iff \mathbb{E}_\mu \left[\min_{p \in \Delta(S)} \mathbb{E}_p[u(f)] + \frac{R(p||q)}{\lambda_q} \right] \geq \mathbb{E}_\mu \left[\min_{p \in \Delta(S)} \mathbb{E}_p[u(g)] + \frac{R(p||q)}{\lambda_q} \right]$$

where $u : X \rightarrow \mathbb{R}$ is a nonconstant affine function, $Q \subseteq \Delta(S)$ is a finite and nonempty set such that there exists an injective map

$$\begin{aligned} Q &\rightarrow \Delta(\Omega) \\ q &\mapsto \rho_q \end{aligned}$$

with

$$q(\{\omega, \rho_q\}) = \rho_q(\omega) \quad \forall q \in Q, \forall \omega \in \Omega,$$

$\mu \in \Delta(Q)$, and $(\lambda_q)_{q \in Q} \in \mathbb{R}_+^Q$. Then:

1. For every $\Omega \times B \in \Sigma_s$ and $f, h \in \mathcal{F}$

$$\begin{aligned} &\int_Q \min_{p \in \Delta(S)} \int_S u(f_{\Omega \times B} h) \, dp + \frac{1}{\lambda_q} R(p||q) \, d\mu(q) \\ &= \int_{\{q \in Q : \rho_q \in B\}} \min_{p \in \Delta(S)} \int_S u(f) \, dp + \frac{1}{\lambda_q} R(p||q) \, d\mu(q) \\ &\quad + \int_{Q \setminus \{q \in Q : \rho_q \in B\}} \min_{p \in \Delta(S)} \int_S u(h) \, dp + \frac{1}{\lambda_q} R(p||q) \, d\mu(q). \end{aligned}$$

2. For every $\Omega \times B \in \Sigma_s$, if $\mu(\{q \in Q : \rho_q \in B\}) = 0$, then $\Omega \times B$ is null.

Proof. 1) Let $\Omega \times B \in \Sigma_s$ and $f, h \in \mathcal{F}$. We have

$$\begin{aligned}
& \int_Q \min_{p \in \Delta(S)} \int_S u(f_{\Omega \times B} h) dp + \frac{1}{\lambda_q} R(p||q) d\mu(q) \\
= & \int_{\{q \in Q: \rho_q \in B\}} \min_{p \in \Delta(S)} \int_S u(f_{\Omega \times B} h) dp + \frac{1}{\lambda_q} R(p||q) d\mu(q) \\
& + \int_{Q \setminus \{q \in Q: \rho_q \in B\}} \min_{p \in \Delta(S)} \int_S u(f_{\Omega \times B} h) dp + \frac{1}{\lambda_q} R(p||q) d\mu(q) \\
= & \int_{\{q \in Q: \rho_q \in B\}} \min_{p \in \Delta(S): q \gg p} \int_S u(f_{\Omega \times B} h) dp + \frac{1}{\lambda_q} R(p||q) d\mu(q) \\
& + \int_{Q \setminus \{q \in Q: \rho_q \in B\}} \min_{p \in \Delta(S): q \gg p} \int_S u(f_{\Omega \times B} h) dp + \frac{1}{\lambda_q} R(p||q) d\mu(q) \\
= & \int_{\{q \in Q: \rho_q \in B\}} \min_{p \in \Delta(S): q \gg p} \int_S u(f) dp + \frac{1}{\lambda_q} R(p||q) d\mu(q) \\
& + \int_{Q \setminus \{q \in Q: \rho_q \in B\}} \min_{p \in \Delta(S): q \gg p} \int_S u(h) dp + \frac{1}{\lambda_q} R(p||q) d\mu(q) \\
= & \int_{\{q \in Q: \rho_q \in B\}} \min_{p \in \Delta(S)} \int_S u(f) dp + \frac{1}{\lambda_q} R(p||q) d\mu(q) \\
& + \int_{Q \setminus \{q \in Q: \rho_q \in B\}} \min_{p \in \Delta(S)} \int_S u(h) dp + \frac{1}{\lambda_q} R(p||q) d\mu(q).
\end{aligned}$$

2) It immediately follows from 1). ■

Lemma 17. Suppose that the assumptions of Theorem 3 hold. Let \succsim be such that for all $f, g \in \mathcal{F}$

$$f \succsim g \iff \mathbb{E}_\mu \left[\min_{p \in \Delta(S)} \mathbb{E}_p[u(f)] + \frac{R(p||q)}{\lambda_q} \right] \geq \mathbb{E}_\mu \left[\min_{p \in \Delta(S)} \mathbb{E}_p[u(g)] + \frac{R(p||q)}{\lambda_q} \right]$$

where $u : X \rightarrow \mathbb{R}$ is a nonconstant affine function, $Q \subseteq \Delta(S)$ is finite, nonempty, and such that there exists an injective map

$$\begin{aligned}
Q & \rightarrow \Delta(\Omega) \\
q & \mapsto \rho_q
\end{aligned}$$

with

$$q(\{\omega, \rho_q\}) = \rho_q(\omega) \quad \forall q \in Q, \forall \omega \in \Omega,$$

$\mu \in \Delta(Q)$, and $(\lambda_q)_{q \in Q} \in \mathbb{R}_+^Q$. Then \succsim satisfies Uniform Misspecification Concern if and only if there exists λ^* with $\lambda_q = \lambda^*$ for all $q \in \text{supp}\mu$.

Proof. (If) Let $\rho, \rho' \in \Delta(\Omega)$, $f, g \in \mathcal{F}$, and $x \in X$ be such that $\Omega \times \{\rho\}$ and $\Omega \times \{\rho'\}$ are nonnull,

$$\rho(\{\omega : f(\omega, \rho) = y\}) = \rho'(\{\omega : g(\omega, \rho') = y\}) \quad \forall y \in X, \quad (24)$$

and $f \succsim_{\Omega \times \{\rho\}} x$. Since $\Omega \times \{\rho\}$ and $\Omega \times \{\rho'\}$ are nonnull, by part 2 of Lemma 16 there exist $q, q' \in Q$ with $\mu(\{q\}) > 0$, $\mu(\{q'\}) > 0$, $\rho_q = \rho$, and $\rho_{q'} = \rho'$. Let

$$\phi(c) = -\exp(-\lambda^* c), \quad \forall c \in u(X)$$

and let $\xi \in \Delta(X)$ be the finite support probability measure such that for all $y \in X$, $\xi(y) = q(\{(\omega, \rho_q) : f(\omega, \rho_q) = y\})$, then

$$\int_{\Omega} \phi(u(f)) \, dq = \int_X \phi(u(y)) \, d\xi(y).$$

Moreover, equation (24) implies

$$\int_{\Omega} \phi(u(g)) \, dq' = \int_X \phi(u(y)) \, d\xi(y).$$

Therefore, by Lemma 16 both $f \succsim_{\Omega \times \{\rho\}} x$ and $g \succsim_{\Omega \times \{\rho'\}} x$ mean that

$$\int_X \phi(u(y)) \, d\xi(y) \geq \phi(u(x))$$

proving that \succsim satisfies Uniform Misspecification Concern.

(Only if) Suppose by way of contradiction that there exist $q, q' \subseteq Q$ and $k \in \mathbb{R}_{++}$ with $\mu(\{q\}) > 0$ and $\mu(\{q'\}) > 0$ and

$$\lambda_q > k > \lambda_{q'}. \quad (25)$$

Since the state space is Q -adequate there exist two events $(W_q, W_{q'}) \subseteq \Delta(\Omega) \times \Delta(\Omega)$ and $c \in (0, 1)$ with

$$\rho_q(W_q) = \rho_{q'}(W_{q'}) = c.$$

Moreover, $q(W_q \times \{\rho_q\}) = c = q'(W_{q'} \times \{\rho_{q'}\})$ and

$$q(W_{q'} \times \{\rho_{q'}\}) = 0 = q'(W_q \times \{\rho_q\}).$$

Pick $z, y \in X$ with $z \succ y$. We have that

$$\begin{aligned} & \rho_q(\{\omega : z((W_q \times \{\rho_q\}) \cup (W_{q'} \times \{\rho_{q'}\})) y(\omega, \rho_q) = x\}) \\ &= \rho_{q'}(\{\omega : z((W_q \times \{\rho_q\}) \cup (W_{q'} \times \{\rho_{q'}\})) y(\omega, \rho_{q'}) = x\}) \end{aligned}$$

for all $x \in X$. By the convexity of X and Lemma 16 there exists $\hat{x} \in X$ with $z \succ \hat{x} \succ y$ and

$$z((W_q \times \{\rho_q\}) \cup (W_{q'} \times \{\rho_{q'}\})) y \sim_{\rho_{q'}} \hat{x}.$$

But by equation (25) and Lemma 16 we have

$$\hat{x} \succ_{\rho_q} z((W_q \times \{\rho_q\}) \cup (W_{q'} \times \{\rho_{q'}\})) y$$

a violation of Uniform Misspecification Concern. ■

Proof of Theorem 3. (Only if) That \succsim satisfies Weak Order, Weak Certainty Independence, Continuity, Monotonicity, Uncertainty Aversion, Nondegeneracy, and Weak Monotone Continuity follows from Lemma 12.

Let $\rho \in \Delta(\Omega)$, $W \subseteq \Omega$, $f, g, h, h' \in \mathcal{F}$, and $fWh \succsim_\rho gWh$. If $\Omega \times \{\rho\}$ is null then we trivially have $fWh' \succsim_\rho gWh'$. Therefore, suppose $\Omega \times \{\rho\}$ is nonnull. By Lemma 16, and since $q \mapsto \rho_q$ is injective, there exists $\bar{q} \in \Delta(S)$ with $\rho_{\bar{q}} = \rho$, $\mu(\{\bar{q}\}) > 0$, and

$$\min_{p \in \Delta(S)} \int_S u(fWh) dp + \frac{1}{\lambda} R(p||\bar{q}) \geq \min_{p \in \Delta(S)} \int_S u(gWh) dp + \frac{1}{\lambda} R(p||\bar{q}).$$

By Proposition 1.4.2 in Dupuis and Ellis (2011) this is equivalent to

$$\int_S \phi(u(fWh)) d\bar{q} \geq \int_S \phi(u(gWh)) d\bar{q}$$

with $\phi(\cdot) = -\exp(-\lambda(\cdot))$. This is also equivalent to

$$\begin{aligned} & \int_{W \times \Delta(\Omega)} \phi(u(f)) d\bar{q} + \int_{(\Omega \setminus W) \times \Delta(\Omega)} \phi(u(h)) d\bar{q} \\ \geq & \int_{W \times \Delta(\Omega)} \phi(u(g)) d\bar{q} + \int_{(\Omega \setminus W) \times \Delta(\Omega)} \phi(u(h)) d\bar{q} \end{aligned} \quad (26)$$

or

$$\int_{W \times \Delta(\Omega)} \phi(u(f)) d\bar{q} \geq \int_{W \times \Delta(\Omega)} \phi(u(g)) d\bar{q}.$$

But then, by reversing all the steps with h' in place of h we get

$$fWh' \succsim_{\rho} gWh'$$

and therefore \succsim satisfies Intramodel Sure-Thing Principle.

Moreover, \succsim satisfies Uniform Misspecification Concern by Lemma 17. That there is a finite set $B \subseteq \Delta(\Omega)$ such that $\Omega \times (\Delta(\Omega) \setminus B)$ is null immediately follows from the representation and part 2 of Lemma 16. Let $\Omega \times B \in \Sigma_s$ and $f, g, h, h' \in \mathcal{F}$. If $\Omega \times B$ is null, we clearly have that $\Omega \times B \in \Sigma_{st}$. Suppose $\Omega \times B$ is nonnull, then

$$\begin{aligned} f(\Omega \times B)h & \succsim g(\Omega \times B)h \\ & \iff \\ \int_{\{q \in Q: \rho_q \in B\}} \min_{p \in \Delta(S)} \mathbb{E}_p[u(f)] + \frac{R(p||q)}{\lambda_q} d\mu(q) & \geq \int_{\{q \in Q: \rho_q \in B\}} \min_{p \in \Delta(S)} \mathbb{E}_p[u(f)] + \frac{R(p||q)}{\lambda_q} d\mu(q) \\ & \iff \\ f(\Omega \times B)h' & \succsim g(\Omega \times B)h' \end{aligned}$$

where the two equivalences follow by Lemma 16. This proves that $\Omega \times B \in \Sigma_{st}$. Since B was chosen to be an arbitrary measurable subset of $\Delta(\Omega)$, $\Sigma_s \subseteq \Sigma_{st}$, and Structured Savage P2 holds.

That \succsim satisfies Structured Savage P4 and Uncertainty Neutrality over Models immediately follows from Lemma 16 and the representation.

(If) By Structured Savage's P2, $\Sigma_s \subseteq \Sigma_{st}$. Suppose $E \in \Sigma_s$ is nonnull, and let $x, x' \in X$ with $x \succ x'$. Then there exist $f, g, h \in \mathcal{F}$ such that $fEh \succ gEh$. Since f and g are simple acts, they assume finitely many values, and by Weak Order, there exist $\bar{x}, \underline{x} \in X$ with

$$\bar{x} \succsim f(s), \quad g(s) \succsim \underline{x}, \quad \forall s \in E.$$

Since $E \in \Sigma_s \subseteq \Sigma_{st}$, $fE\bar{x} \succ gE\bar{x}$. By the Monotonicity and Weak Order parts of the Variational Axiom, $\underline{x}\emptyset\bar{x} = \bar{x}E\bar{x} \succsim fE\bar{x} \succ gE\bar{x} \succsim \underline{x}E\bar{x}$. Therefore, by Structured Savage P4, $x = x'\emptyset x \succ x'Ex$. Since $E \in \Sigma_s$ and $x, x' \in X$ were arbitrarily chosen, each nonnull $E \in \Sigma_s$ is also strongly nonnull.

Next, fix a finite B , such that for each $\rho \in B$, $\Omega \times \{\rho\}$ is nonnull, and such that $S \setminus \{\Omega \times B\}$ is null. Such a set exists by the Structured Savage axiom. For every $\rho \in B$, by the previous part of the proof $\Omega \times \{\rho\}$ is strongly nonnull and so by Lemma 14 we have

$$f \succsim_\rho g \iff \min_{p \in \Delta(S)} \int_S \hat{u}(f) dp + \frac{1}{\lambda_\rho} R(p||q_\rho) \quad (27)$$

for some $q_\rho \in \Delta(S)$ with support contained in $\Omega \times \{\rho\}$ and a nonconstant affine \hat{u} . In particular, since the space is Q -adequate, by applying Uniform Misspecification Concern with $\rho = \rho'$, we obtain that $q_\rho = \rho \times \delta_\rho$. Let

$$Q = \{q_\rho \in \Delta(S) : \rho \in B\}. \quad (28)$$

Identify each act $f \in \mathcal{F}$ with the real-valued function $\hat{f} : Q \rightarrow \hat{u}(X)$ with

$$\hat{f}(q) = \min_{p \in \Delta(S)} \int_S \hat{u}(f) dp + \frac{1}{\lambda_{\rho_q}} R(p||q_{\rho_q}) \quad \forall q \in Q$$

where λ_{ρ_q} is given by equation (27).

We now show that

$$\hat{f} = \hat{g} \implies f \sim g \quad \forall f, g \in \mathcal{F}.$$

We partition S in $\left\{ \{\Omega \times \rho\}_{\rho \in B}, S \setminus \{\Omega \times B\} \right\}$ and establish the claim by induction on the number of cells of the partition on which f and g are not identical. Let f and g be such that $\hat{f} = \hat{g}$ and they differ on one element of the partition, say E . Then $f = fEg \sim g$ by definition of \sim_E and Structured Savage P2, so $f \sim g$. For the inductive step, suppose that whenever f and g are such that $\hat{f} = \hat{g}$ and they differ at most on $n \in \mathbb{N}$ elements of the partition, we have $f \sim g$. Let f and g be such that $\hat{f} = \hat{g}$ and they differ on $n+1 \in \mathbb{N}$ elements of the partition. Let E be an element of the partition on which they differ. Then, fEg and g differ on one element of the partition, and fEg and f differ on n elements of the partition. Therefore, by the inductive hypothesis, we have $g \sim fEg \sim f$.

Moreover, it is immediate to see that $\hat{u}(X)^Q \subseteq \left\{ \hat{f} : f \in \mathcal{F} \right\}$. Therefore, with a slight abuse of notation we let \succsim denote also the binary relation on $\hat{u}(X)^Q$ defined

by $\hat{f} \succsim \hat{g}$ if and only if $f \succsim g$. With this, by the Continuity part of the Variational Axiom, Structured Savage, Uncertainty Neutrality over Models, and Theorem VII.3.5 in Wakker (2013) there exists $\mu \in \Delta(Q)$ such that for all $\psi, \psi' \in \hat{u}(X)^Q$

$$\psi \succsim \psi' \iff \sum_{q \in Q} \psi(q) \mu(q) \geq \sum_{q \in Q} \psi'(q) \mu(q).$$

Moreover, it is immediate to see that μ is uniquely identified.

But then, by definition of \succsim , we obtain that for all $f, g \in \mathcal{F}$

$$f \succsim g \iff \mathbb{E}_\mu \left[\min_{p \in \Delta(S)} \mathbb{E}_p[u(f)] + \frac{R(p||q)}{\lambda_q} \right] \geq \mathbb{E}_\mu \left[\min_{p \in \Delta(S)} \mathbb{E}_p[u(g)] + \frac{R(p||q)}{\lambda_q} \right].$$

Moreover, by Lemma 17 Uniform Misspecification Concern implies that $\lambda = \lambda_E$ for all $E \in \Sigma_s$, proving the result. \blacksquare

Proof of Corollary 2. It immediately follows from Lemma 12 and Proposition 8 in Maccheroni, Marinacci, and Rustichini (2006a). \blacksquare

Proof of Proposition 4. That each \succsim^h admits an average robust control representation $(u_h, Q_h, \mu_h, \lambda_h)$ follows from Theorem 3. That $u_h = u$ for some constant affine u follows from Constant Preference Invariance, and λ_h is increasing in $LRT(h, Q)$ by Proposition 8 in Maccheroni, Marinacci, and Rustichini (2006a) and Q -Likelihood.

We now prove that Dynamic Consistency over Models implies $\mu(\cdot|h_t) = \mu_{h_t}$ for all $h_t = (\omega_i)_{i=1}^t \in \mathcal{H}_t$ such that $\prod_{i=1}^t \rho_q(\omega_i) > 0$ for some $q \in Q$. By definition, we have $\mu_{h_t} = \mu$ for the empty history. Let f and g be measurable with respect to Σ_s . Then we can suppress the dependence on ω in $f(\omega, \rho)$ and $g(\omega, \rho)$ and we have that

$$f \succsim^{h_t} g \iff f^0 \succsim^\emptyset g^0.$$

But by construction, the latter is equivalent to

$$\begin{aligned} & \int_{\Delta(\Delta(S))} \gamma_{f(\rho_q)} \prod_{i=1}^t \rho_q(\omega_i) (u(\bar{z}) - u(\underline{z})) d\mu(q) \\ & \geq \int_{\Delta(\Delta(S))} \gamma_{g(\rho_q)} \prod_{i=1}^t \rho_q(\omega_i) (u(\bar{z}) - u(\underline{z})) d\mu(q). \end{aligned}$$

Dividing both side by the strictly positive ex-ante probability of history h_t , we obtain

$$\begin{aligned} & \frac{\int_{\Delta(\Delta(S))} \gamma_{f(\rho_q)} \prod_{i=1}^t \rho_q(\omega_i) (u(\bar{z}) - u(\underline{z})) d\mu(q)}{\int_{\Delta(\Delta(S))} \prod_{i=1}^t \rho_q(\omega_i) d\mu(q)} \\ & \geq \frac{\int_{\Delta(\Delta(S))} \gamma_{g(\rho_q)} \prod_{i=1}^t \rho_q(\omega_i) (u(\bar{z}) - u(\underline{z})) d\mu(q)}{\int_{\Delta(\Delta(S))} \prod_{i=1}^t \rho_q(\omega_i) d\mu(q|h_t)}. \end{aligned}$$

But then, by the formula for Bayesian updating, this is equivalent to

$$\int_{\Delta(\Delta(S))} \gamma_{f(\rho_q)} (u(\bar{z}) - u(\underline{z})) d\mu(q|h_t) \geq \int_{\Delta(\Delta(S))} \gamma_{g(\rho_q)} (u(\bar{z}) - u(\underline{z})) d\mu(q|h_t)$$

that is

$$\int_{\Delta(\Delta(S))} u(f(\rho_q)) d\mu(q|h_t) \geq \int_{\Delta(\Delta(S))} u(g(\rho_q)) d\mu(q|h_t).$$

That is, \succsim^{h_t} admits an SEU representation of the acts measurable with respect to Σ_s with Bernoulli utility u and probability measure $\mu(\cdot|h_t)$. Since for the histories $h_t = (\omega_i)_{i=1}^t \in \mathcal{H}_t$ where $\prod_{i=1}^t \rho_q(\omega_i) = 0$ for all $q \in Q$ Bayesian updating does not impose any restriction, the result follows. \blacksquare

Proof of Proposition 5. By Proposition 4, \succsim^h admits an average robust control representation $(u, Q, \mu(\cdot|h), \lambda_h)$ for every $h \in \mathcal{H}$. Observe that since the outcome frequency is constant along the sequence $(h_{t_n})_{n \in \mathbb{N}}$, by Lemma 2, $\frac{LRT(h_{t_n}, Q)}{2t_n} = 1/c$ for some $c \in \mathbb{R}_{++}$ and for all $n \in \mathbb{N}$. Let $\rho \in \Delta(\Omega)$. Suppose by way of contradiction that

$$\begin{aligned} l: &= \liminf_{n \rightarrow \infty} c\lambda^{h_{t_n}} = \liminf_{n \rightarrow \infty} \lambda_{h_{t_n}} / \left(\frac{LRT(h_{t_n}, Q)}{2t_n} \right) \\ &< \limsup_{n \rightarrow \infty} \lambda_{h_{t_n}} / \left(\frac{LRT(h_{t_n}, Q)}{2t_n} \right) = \limsup_{n \rightarrow \infty} c\lambda_{h_{t_n}} =: L. \end{aligned}$$

Let $\bar{q} \in Q$ be such that $\Omega \times \{\rho_{\bar{q}}\}$ is nonnull. Since $\Omega \times \{\rho_{\bar{q}}\}$ contains at least three nonnull events, there is E and $r \in (0, 1)$ with $\rho_{\bar{q}}(E) = r$. Let $x, z \in X$, $\gamma^*, \gamma_* \in (0, 1)$, and $\lambda^*, \lambda_* \in (\frac{l}{c}, \frac{L}{c})$ be such that $x \succ^\emptyset z$, $\lambda^* > \lambda_*$,

$$\begin{aligned} & -\mu(q^*) \frac{\log(r \exp(-\lambda^*(u(z))) + (1-r) \exp(-\lambda^*(u(x))))}{\lambda^*} \\ & + (1 - \mu(q^*)) u(\gamma^*x + (1 - \gamma^*)z) \\ & = u(\gamma^*x + (1 - \gamma^*)z), \end{aligned}$$

and

$$\begin{aligned}
& -\mu(q^*) \frac{\log(r \exp(-\lambda_*(u(z))) + (1-r) \exp(-\lambda_*(u(x))))}{\lambda_*} \\
& + (1 - \mu(q^*)) u(\gamma_* x + (1 - \gamma_*) z) \\
= & u(\gamma_* x + (1 - \gamma_*) z),
\end{aligned}$$

where the existence of such γ_*, γ^* is guaranteed by u being affine. Moreover, it is easy to see that $\gamma_* > \gamma^*$. Consider a subsequence $(n_m)_{m \in \mathbb{N}}$ such that

$$\lim_{m \rightarrow \infty} c\lambda^{h_{t_{n_m}}} = l.$$

Moreover, let $M \in \mathbb{N}$ be such that for all $m \geq M$

$$\lambda^{h_{t_{n_m}}} < \frac{\lambda_* + \frac{l}{c}}{2}.$$

Similarly, let $(n_{\tilde{m}})_{\tilde{m} \in \mathbb{N}}$ such that

$$\lim_{\tilde{m} \rightarrow \infty} c\lambda^{h_{t_{n_{\tilde{m}}}}} = L.$$

Moreover, let $\tilde{M} \in \mathbb{N}$ be such that for all $\tilde{m} \geq \tilde{M}$

$$\lambda^{h_{t_{n_{\tilde{m}}}}} > \frac{\lambda^* + \frac{L}{c}}{2}.$$

With this, by Proposition 4 and Lemma 15 we have that for all $m \geq M$ and $\tilde{m} \geq \tilde{M}$

$$\gamma_{\succsim_{h_{t_{n_m}}}}^{xEz} > \gamma_* \text{ and } \gamma_{\succsim_{h_{t_{n_{\tilde{m}}}}}}^{xEz} < \gamma^*.$$

But this in turn implies that $\succsim_{h_{t_{n_m}}}$ is never $(x, y, E, (\gamma_* - \gamma^*))$ -similar to $\succsim_{h_{t_{n_{\tilde{m}}}}}$ for

$$\min\{m, \tilde{m}\} \geq \max\{M, \tilde{M}\},$$

a contradiction. This shows that either $\lambda_{h_{t_n}}$ converges or it diverges to plus infinity. The last part of the statement immediately by taking E in the first part of the proof to be equal to the one whose existence is asserted in the statement, and by the construction of γ_* and γ^* above. ■

Proof of Proposition 6. By Proposition 4 we know that each \succsim^h admits an average robust control representation $(u, Q, \mu(\cdot|h), \lambda_h)$. Observe that since the outcome frequency is constant along the sequence $(h_{t_n})_{n \in \mathbb{N}}$, by Lemma 2, $\frac{LRT(h_{t_n}, Q)}{t_n} = c$ for some $c \in \mathbb{R}_{++}$ and for all $n \in \mathbb{N}$. Suppose by way of contradiction that

$$L := \limsup_{n \rightarrow \infty} \frac{LRT(h_{t_n}, Q)}{\lambda_{h_{t_n}} t_n} > 0 \quad (29)$$

for some sequence of histories with empirical frequency $\hat{\rho} \notin \{\rho_q : q \in Q\}$. Since the state space is Q -adequate there exist

$$k \in (0, 1) \text{ and } (W_q)_{q \in Q} \in (2^\Omega)^Q$$

such that $\rho_q(W_q) = k$ for all $q \in Q$. Define $E = \cup_{q \in Q} (\{\rho_q\} \times W_q)$. Let $x, y \in X$ with $x \succ^\emptyset y$ and $f \in \mathcal{F}$ be defined as

$$f = xEy.$$

Choose also $z \in X$ such that $x \succ^\emptyset z \succ^\emptyset y$ and

$$-\exp\left(-2\frac{c}{L}u(z)\right) = -k \exp\left(-2\frac{c}{L}u(x)\right) - (1-k) \exp\left(-2\frac{c}{L}u(y)\right)$$

where the existence of such z is guaranteed by u being affine and X being convex. But then equation (29) implies that for infinitely many $n \in \mathbb{N}$

$$f \succ_\rho^{h_{t_n}} z$$

a contradiction with Asymptotic Concern for every $\rho \in \{\rho_q : q \in Q\}$. ■

Proof of Proposition 7. By Proposition 4 we know that each \succsim^h admits an average robust control representation $(u, Q, \mu(\cdot|h), \lambda_h)$. Consider a sequence of histories $(h_{t_n})_{n \in \mathbb{N}}$ with outcome frequency constant and not in $\{\rho_q : q \in Q\}$. Observe that since the outcome frequency is constant along the sequence $(h_{t_n})_{n \in \mathbb{N}}$, by Lemma 2, $\frac{LRT(h_{t_n}, Q)}{t_n} = c$ for some $c \in \mathbb{R}_{++}$ and for all $n \in \mathbb{N}$. Suppose by way of contradiction that

$$L := \liminf_{n \rightarrow \infty} \frac{LRT(h_{t_n}, Q)}{\lambda_{h_{t_n}} t_n} \in \mathbb{R}. \quad (30)$$

As the state space is Q -adequate there exist $k \in (0, 1)$ and $(W_q)_{q \in Q} \in (2^\Omega)^Q$ such that

for every $q \in Q$, $\rho_q(W_q) = k$. Let $x, z \in X$ and $\gamma \in (0, 1)$ be such that $x \succ^\emptyset z$ and

$$\begin{aligned} & -k \exp \left(-\frac{c}{2 \max \{1, L\}} (u(x)) \right) - (1-k) \exp \left(-\frac{c}{2 \max \{1, L\}} (u(z)) \right) \\ &= -\exp \left(-\frac{c}{2 \max \{1, L\}} (u(\gamma x + (1-\gamma)z)) \right), \end{aligned}$$

where the existence of such γ is guaranteed by u being affine, and $\gamma < k$. Define $E = \cup_{q \in Q} (W_q \times \{\rho_q\})$. Consider a subsequence $(n_m)_{m \in \mathbb{N}}$ such that

$$\lim_{m \rightarrow \infty} \frac{LRT(h_{t_{n_m}}, Q)}{\lambda_{h_{t_{n_m}}} t_{n_m}} = L.$$

Moreover, let M be such that for all $m \geq M$

$$\lambda^{h_{t_{n_m}}} / c > \frac{1}{2 \max \{1, L\}}.$$

With this, by Proposition 4 if we let \geq be the subjective utility preference with utility index u and belief $\int_Q p d\mu(q)$, we have that for all $m \geq M$

$$\gamma_{\sim^{h_{t_{n_m}}}}^{xEz} > \gamma \text{ and } \gamma_{\geq^{h_{t_{n_m}}}}^{xEz} = k.$$

By Corollary 2, this contradicts Asymptotic Leniency as then $\sim^{h_{t_{n_m}}}$ and $\geq^{h_{t_{n_m}}}$ are not $(x, y, E, k - \gamma)$ -similar for $m \geq M$. ■

A.3 General Statistical Distances

The results of the paper that involve the statistically sophisticated type extend easily to the case of an average of general divergence preferences (Cerreia-Vioglio, Hansen, Maccheroni, and Marinacci, 2022), i.e., to decision criteria of the form

$$\int_{\Theta} \min_{p \in \Delta(S)} \mathbb{E}_{p_a} [u(a, y)] + \frac{1}{\lambda} D_{\phi}(p_a || q_a^{\theta}) d\mu(\theta)$$

where

$$D_{\phi}(p_a || q_a^{\theta}) = \begin{cases} \int \phi \left(\frac{dp_a}{dq_a} \right) dq_a^{\theta} & q_a^{\theta} \gg p_a \\ \infty & \text{otherwise} \end{cases}$$

for some continuous strictly convex function $\phi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ with

$$\phi(1) = 0 \text{ and } \lim_{t \rightarrow \infty} \frac{\phi(t)}{t} = \infty.$$

From this expression, it is clear that the main case studied in the paper is the one where $\phi(c) = c \log c - c + 1$. The only caveat is that the best reply function BR^λ must now be defined with respect to the relevant divergence.

A.4 Computations supporting

A.4.1 Example 1

We have

$$\begin{aligned} R(p_a^* || q_a^\theta) &= \text{const.} \\ &+ \int_{\mathbb{R}} \int_{\mathbb{R}} \frac{-\exp\left(-\frac{(t-\tau(a+\omega_a))^2}{2}\right)}{\sqrt{2\pi}} \log\left(\frac{\exp\left(-\left(\frac{t}{a+\omega_a} - \theta\right)^2/2\right)}{\sqrt{2\pi}}\right) dt dp_a^*(\omega_a) \\ &= \text{const.} + \int_{\mathbb{R}} \int_{\mathbb{R}} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{(t-\tau(a+\omega_a))^2}{2}\right) \left(\frac{t}{a+\omega_a} - \theta\right)^2 / 2 dt dp_a^*(\omega_a) \end{aligned}$$

proving that $\Theta(a) = \mathbb{E}_{p_a^*} \left[\frac{\tau(a+\omega_a)}{a+\omega_a} \right]$. The condition for not switching from an action a to an action a' with $a \geq a'$ in a Berk-Nash equilibrium in which the belief is concentrated on θ is

$$\mathbb{E}_{\omega_a, \varepsilon_2} [(a + \omega_a)(1 - \theta - \varepsilon_2)] - \mathbb{E}_{\omega_{a'}, \varepsilon_2} [(a' + \omega_{a'})(1 - \theta - \varepsilon_2)] \geq c(a) - c(a').$$

By Proposition 1.4.2 in Dupuis and Ellis (2011), the condition for not switching from an action a to an action a' with $a \geq a'$ in a k -robust equilibrium in which the

belief is concentrated on θ is

$$\begin{aligned}
& -\frac{k}{R(p_a^*||q_a^\theta)} \log \mathbb{E}_{\omega_a, \omega_{a'}, \varepsilon_2} \left[\exp \left(-\frac{R(p_a^*||q_a^\theta) [(a + \omega_a)(1 - \theta - \varepsilon_2) - (a' + \omega_{a'})(1 - \theta - \varepsilon_2)]}{k} \right) \right] \\
& = -\frac{k}{R(p_a^*||q_a^\theta)} \log \mathbb{E}_{\omega_a, \omega_{a'}, \varepsilon_2} \left[\exp \left(-\frac{R(p_a^*||q_a^\theta) (a + \omega_a)(1 - \theta - \varepsilon_2)}{k} \right) \right] \\
& \quad + \frac{k}{R(p_a^*||q_a^\theta)} \log \mathbb{E}_{\omega_a, \omega_{a'}, \varepsilon_2} \left[\exp \left(-\frac{R(p_a^*||q_a^\theta) (a' + \omega_{a'})(1 - \theta - \varepsilon_2)}{k} \right) \right] \\
& \geq c(a) - c(a').
\end{aligned}$$

Since the LHS is lower in the second case, we obtain the desired conclusion.

A.4.2 Example 2

The condition for not switching from action 0 to an action a with $a \geq 0$ in a Berk-Nash equilibrium in which the belief is concentrated on θ is

$$p_a^*(s \leq a) (\mathbb{E}(v) - a) \leq 0.$$

By Proposition 1.4.2 in Dupuis and Ellis (2011), the condition for not switching from action 0 to an action a' with $a \geq a'$ in a k -robust equilibrium in which the belief is concentrated on θ is

$$-\frac{k}{R(p_a^*||q_a^\theta)} \log \int_{\mathbb{R}} \int_{\mathbb{R}} \exp \left(-\frac{R(p_a^*||q_a^\theta)}{k} [v - a] \mathbb{I}_{[0, a]}(s) \right) dp_a^*(s) dp_a^*(v) \leq 0.$$

Since the LHS is lower in the second case, we obtain the desired conclusion.

References

- Aliprantis, C. and K. Border (2013). *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer-Verlag.
- Anderson, R. M., H. Duanmu, A. Ghosh, and M. A. Khan (2022). “On Existence of Berk-Nash Equilibria in Misspecified Markov Decision Processes with Infinite Spaces”. *arXiv preprint arXiv:2206.08437*.

- Angeletos, G.-M., Z. Huo, and K. A. Sastry (2021). “Imperfect macroeconomic expectations: Evidence and theory”. *NBER Macroeconomics Annual* 35, pp. 1–86.
- Arrow, K. and J. Green (1973). “Notes on Expectations Equilibria in Bayesian Settings”. Working Paper No. 33, Stanford University.
- Aubin, J.-P. and A. Cellina (2012). *Differential inclusions: set-valued maps and viability theory*. Vol. 264. Springer Science & Business Media.
- Ba, C. (2022). “Robust Model Misspecification and Paradigm Shifts”. *arXiv preprint arXiv:2106.12727*.
- Banerjee, A. V., S. Chassang, S. Montero, and E. Snowberg (2020). “A theory of experimenters: Robustness, randomization, and balance”. *American Economic Review* 110, pp. 1206–30.
- Battigalli, P. (1987). “Comportamento razionale ed equilibrio nei giochi e nelle situazioni sociali”. *Unpublished undergraduate dissertation, Bocconi University, Milano*.
- Battigalli, P., S. Cerreia-Vioglio, F. Maccheroni, M. Marinacci, and T. Sargent (2022). “A framework for the analysis of self-confirming policies”. *Theory and Decision*, pp. 1–58.
- Battigalli, P., A. Francetich, G. Lanzani, and M. Marinacci (2019). “Learning and self-confirming long-run biases”. *Journal of Economic Theory* 183, pp. 740–785.
- Benaim, M., J. Hofbauer, and S. Sorin (2005). “Stochastic approximations and differential inclusions”. *SIAM Journal on Control and Optimization* 44, pp. 328–348.
- Berk, R. H. (1966). “Limiting behavior of posterior distributions when the model is incorrect”. *The Annals of Mathematical Statistics*, pp. 51–58.
- Billingsley, P. (2008). *Probability and measure*. John Wiley & Sons.
- Bohren, J. A. and D. N. Hauser (2021). “Learning with heterogeneous misspecified models: Characterization and robustness”. *Econometrica* 89, pp. 3025–3077.
- Cabrales, A., O. Gossner, and R. Serrano (2013). “Entropy and the value of information for investors”. *American Economic Review* 103, pp. 360–77.
- Casella, G. and R. L. Berger (2021). *Statistical inference*. Cengage Learning.
- Cerreia-Vioglio, S., L. P. Hansen, F. Maccheroni, and M. Marinacci (2022). “Making decisions under model misspecification”.
- Cerreia-Vioglio, S., F. Maccheroni, M. Marinacci, and L. Montrucchio (2011). “Uncertainty averse preferences”. *Journal of Economic Theory* 146, pp. 1275–1330.
- (2013a). “Ambiguity and robust statistics”. *Journal of Economic Theory* 148, pp. 974–1049.

- Cerreia-Vioglio, S., F. Maccheroni, M. Marinacci, and L. Montrucchio (2013b). “Classical subjective expected utility”. *Proceedings of the National Academy of Sciences* 110, pp. 6754–6759.
- Chew, H. S. and J. S. Sagi (2006). “Event exchangeability: Probabilistic sophistication without continuity or monotonicity”. *Econometrica* 74, pp. 771–786.
- Cho, I.-K. and K. Kasa (2015). “Learning and model validation”. *The Review of Economic Studies* 82, pp. 45–82.
- Clarida, R., J. Gali, and M. Gertler (2000). “Monetary policy rules and macroeconomic stability: evidence and some theory”. *The Quarterly journal of economics* 115, pp. 147–180.
- Cogley, T. and T. J. Sargent (2005). “The conquest of US inflation: Learning and robustness to model uncertainty”. *Review of Economic dynamics* 8, pp. 528–563.
- Dean, M. and P. Ortoleva (2017). “Allais, Ellsberg, and preferences for hedging”. *Theoretical Economics* 12, pp. 377–424.
- (2019). “The empirical relationship between nonstandard economic behaviors”. *Proceedings of the National Academy of Sciences* 116, pp. 16262–16267.
- Dekel, E., J. C. Ely, and O. Yilankaya (2007). “Evolution of preferences”. *The Review of Economic Studies* 74, pp. 685–704.
- Denti, T. and L. Pomatto (2022). “Model and predictive uncertainty: A foundation for smooth ambiguity preferences”. *Econometrica* 90, pp. 551–584.
- Diamond, P. A. (1982). “Aggregate demand management in search equilibrium”. *Journal of political Economy* 90, pp. 881–894.
- Dillenberger, D., J. S. Lleras, P. Sadowski, and N. Takeoka (2014). “A theory of subjective learning”. *Journal of Economic Theory* 153, pp. 287–312.
- Dudley, R. M. (2018). *Real Analysis and Probability*. Chapman and Hall/CRC.
- Dunford, N. and J. T. Schwartz (1988). *Linear operators, part 1: general theory*. Vol. 10. John Wiley & Sons.
- Dupuis, P. and R. S. Ellis (2011). *A weak convergence approach to the theory of large deviations*. Vol. 902. John Wiley & Sons.
- Enke, B. and F. Zimmermann (2019). “Correlation neglect in belief formation”. *The Review of Economic Studies* 86, pp. 313–332.
- Epstein, L. G. and S. Ji (2022). “Optimal learning under robustness and time-consistency”. *Operations Research* 70, pp. 1317–1329.
- Esponda, I. (2008). “Behavioral equilibrium in economies with adverse selection”. *American Economic Review* 98, pp. 1269–91.

- Esponda, I. and D. Pouzo (2016). “Berk–Nash equilibrium: A Framework for Modeling Agents with Misspecified Models”. *Econometrica* 84, pp. 1093–1130.
- Esponda, I., D. Pouzo, and Y. Yamamoto (2021a). “Asymptotic behavior of Bayesian learners with misspecified models”. *Journal of Economic Theory* 195, p. 105260.
- (2021b). “Corrigendum: Asymptotic behavior of Bayesian learners with misspecified models”. *Journal of Economic Theory* 195, p. 105260.
- Esponda, I., E. Vespa, and S. Yuksel (2022). *Mental models and learning: The case of base-rate neglect*. Tech. rep.
- Fishburn, P. C. (1970). *Utility theory for decision making*. Wiley.
- Foster, D. P. and H. P. Young (2003). “Learning, hypothesis testing, and Nash equilibrium”. *Games and Economic Behavior* 45, pp. 73–96.
- Foutz, R. V. and R. Srivastava (1977). “The performance of the likelihood ratio test when the model is incorrect”. *The annals of Statistics*, pp. 1183–1194.
- Frick, M., R. Iijima, and Y. Ishii (2022). “Belief Convergence under Misspecified Learning: A Martingale Approach”. *Review of Economic Studies* forthcoming.
- Frick, M., R. Iijima, and Y. Le Yaouanq (2022). “Objective rationality foundations for (dynamic) α -MEU”. *Journal of Economic Theory* 200, p. 105394.
- Fudenberg, D. and D. M. Kreps (1993). “Learning mixed equilibria”. *Games and economic behavior* 5, pp. 320–367.
- Fudenberg, D. and G. Lanzani (2022). “Which misperceptions persist?” *Theoretical Economics, Forthcoming*.
- Fudenberg, D., G. Lanzani, and P. Strack (2021). “Limit points of endogenous misspecified learning”. *Econometrica* 89, pp. 1065–1098.
- (2022a). “Pathwise Concentration Bounds for Misspecified Bayesian Beliefs”. *Available at SSRN 3805083*.
- (2022b). “Selective Memory Equilibrium”. *Available at SSRN 4015313*.
- Fudenberg, D. and D. K. Levine (1993). “Self-confirming equilibrium”. *Econometrica*, pp. 523–545.
- (1995). “Consistency and cautious fictitious play”. *Journal of Economic Dynamics and Control* 19, pp. 1065–1089.
- Gagnon-Bartsch, T., M. Rabin, and J. Schwartzstein (2022). “Channeled Attention and Stable Errors”.
- Ghirardato, P. and M. Marinacci (2002). “Ambiguity made precise: A comparative foundation”. *Journal of Economic Theory* 102, pp. 251–289.

- Giacomini, R., V. Skreta, and J. Turén (2015). “Models, inattention and expectation updates”.
- Gilboa, I., S. Minardi, and L. Samuelson (2020). “Theories and cases in decisions under uncertainty”. *Games and Economic Behavior* 123, pp. 22–40.
- Gilboa, I. and D. Schmeidler (1989). “Maxmin expected utility with non-unique prior”. *Journal of Mathematical Economics* 18, pp. 141–153.
- Gul, F. and W. Pesendorfer (2014). “Expected uncertain utility theory”. *Econometrica* 82, pp. 1–39.
- Hall, P. and C. C. Heyde (2014). *Martingale limit theory and its application*. Academic press.
- Hansen, L. P. and T. J. Sargent (2001). “Robust control and model uncertainty”. *American Economic Review* 91, pp. 60–66.
- (2005). “Robust estimation and control under commitment”. *Journal of Economic Theory* 124, pp. 258–301.
- (2007). “Recursive robust estimation and control without commitment”. *Journal of Economic Theory* 136, pp. 1–27.
- (2011). “Robustness”. *Robustness*. Princeton university press.
- (2022a). “Risk, Ambiguity, and Misspecification: Decision Theory, Robust Control, and Statistics”.
- (2022b). *Structured Ambiguity and Model Misspecification*.
- Hansen, L. P., T. J. Sargent, G. Turmuhambetova, and N. Williams (2006). “Robust control and model misspecification”. *Journal of Economic Theory* 128, pp. 45–90.
- Hausman, J. A. (1978). “Specification tests in econometrics”. *Econometrica: Journal of the econometric society*, pp. 1251–1271.
- He, K. and J. Libgober (2022). “Evolutionarily stable (mis) specifications: Theory and applications”.
- Kallenberg, O. (1973). “Random measures”. *Random Measures*. De Gruyter.
- Karni, E. and M.-L. Vierø (2013). ““Reverse Bayesianism”: A choice-based theory of growing awareness”. *American Economic Review* 103, pp. 2790–2810.
- Klibanoff, P., M. Marinacci, and S. Mukerji (2005). “A smooth model of decision making under ambiguity”. *Econometrica* 73, pp. 1849–1892.
- (2009). “Recursive smooth ambiguity preferences”. *Journal of Economic Theory* 144, pp. 930–976.
- Kreps, D. (1988). *Notes On The Theory Of Choice*. Westview Press.
- Liebman, J. B. and R. J. Zeckhauser (2004). “Schmeduling”.

- Liese, F. and I. Vajda (1987). *Convex statistical distances*. Vol. 95. Teubner.
- Maccheroni, F., M. Marinacci, and A. Rustichini (2006a). “Ambiguity aversion, robustness, and the variational representation of preferences”. *Econometrica* 74, pp. 1447–1498.
- (2006b). “Dynamic variational preferences”. *Journal of Economic Theory* 128, pp. 4–44.
- Maenhout, P. J. (2004). “Robust portfolio rules and asset pricing”. *Review of financial studies* 17, pp. 951–983.
- Massari, F. and J. Newton (2020). “Learning and equilibrium in misspecified models”. *Available at SSRN 3473630*.
- Mu, X., L. Pomatto, P. Strack, and O. Tamuz (2021). “Monotone additive statistics”. *arXiv preprint arXiv:2102.00618*.
- Nyarko, Y. (1991). “Learning in Mis-specified Models and the Possibility of Cycles”. *Journal of Economic Theory* 55, pp. 416–427.
- Ok, E. A. (2011). “Real analysis with economic applications”. *Real Analysis with Economic Applications*. Princeton University Press.
- Ortoleva, P. (2012). “Modeling the change of paradigm: Non-Bayesian reactions to unexpected news”. *American Economic Review* 102, pp. 2410–36.
- Parthasarathy, K. R. (2005). *Probability Measures on Metric Spaces*. American Mathematical Soc.
- Primiceri, G. E. (2005). “Time varying structural vector autoregressions and monetary policy”. *The Review of Economic Studies* 72, pp. 821–852.
- Rabin, M. (2002). “Inference by believers in the law of small numbers”. *The Quarterly Journal of Economics* 117, pp. 775–816.
- Rees-Jones, A. and D. Taubinsky (2020). “Measuring “schmeduling””. *The Review of Economic Studies* 87, pp. 2399–2438.
- Reny, P. J. (1999). “On the existence of pure and mixed strategy Nash equilibria in discontinuous games”. *Econometrica* 67, pp. 1029–1056.
- Robatto, R. and B. Szentes (2017). “On the biological foundation of risk preferences”. *Journal of Economic Theory* 172, pp. 410–422.
- Royden, H. L. and P. Fitzpatrick (1988). *Real analysis*. Vol. 32. Macmillan New York.
- Rudin, W. (1970). *Real and Complex Analysis*. McGraw-Hill.
- Sargent, T. J. (1999). “The conquest of American inflation”. *The Conquest of American Inflation*. Princeton University Press.

- Sargent, T. J. (2008). “Evolution and intelligent design”. *American Economic Review* 98, pp. 5–37.
- Savage, L. J. (1954). *The foundations of statistics*. Courier Corporation.
- Schwartzstein, J. and A. Sunderam (2021). “Using models to persuade”. *American Economic Review* 111, pp. 276–323.
- Serfozo, R. (1982). “Convergence of Lebesgue integrals with varying measures”. *Sankhyā: The Indian Journal of Statistics, Series A*, pp. 380–402.
- Simon, J. (2020). *Continuous Functions*. John Wiley & Sons.
- Sims, C. A. and T. Zha (2006). “Were there regime switches in US monetary policy?” *American Economic Review* 96, pp. 54–81.
- Siniscalchi, M. (2011). “Dynamic choice under ambiguity”. *Theoretical Economics* 6, pp. 379–421.
- Sobel, J. (1984). “Non-linear prices and price-taking behavior”. *Journal of Economic Behavior & Organization* 5, pp. 387–396.
- Strzalecki, T. (2011). “Axiomatic foundations of multiplier preferences”. *Econometrica* 79, pp. 47–73.
- Sweeting, T. (1986). “On a converse to Scheffé’s theorem”. *The Annals of Statistics* 14, pp. 1252–1256.
- Tversky, A. and D. Kahneman (1971). “Belief in the law of small numbers.” *Psychological bulletin* 76, p. 105.
- Van der Vaart, A. and J. A. Wellner (1997). *Weak convergence and empirical processes with applications to statistics*. Springer.
- Vuong, Q. H. (1989). “Likelihood ratio tests for model selection and non-nested hypotheses”. *Econometrica*, pp. 307–333.
- Wakker, P. P. (2013). *Additive representations of preferences: A new foundation of decision analysis*. Vol. 4. Springer Science & Business Media.
- Willard, S. (2012). *General topology*. Courier Corporation.