Dynamic Concern for Misspecification*

Giacomo Lanzani - MIT

October 2022 - 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 preferences under uncertainty can 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 rationalizes the empirical evidence on monetary policy cycles and labor supply 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, Yeon-Koo Che, Daniel Clark, Tommaso Denti, Esther Duflo, Federico Echenique, Andrew Ellis, Glenn Ellison, Joel Flynn, Mira Frick, Benjamin Golub, Johannes Hörner, Ryota Iijima, Peter Klibanoff, Jonathan Libgober, Fabio Maccheroni, Massimo Marinacci, Moritz Meyer-ter-Vehn, Anh Nguyen, Daisuke Oyama, Luciano Pomatto, Frank Schilbach, Advik Shreekumar, Charles Sprenger, Tomasz Strzalecki, Omer Tamuz, Jean Tirole, Emanuel Vespa, and Alexander Wolitzky for helpful comments and suggestions. The financial support of Guido Cazzavillan 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.[1]

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. 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 full-fledged decision criterion that accounts for model misspecification was proposed in Hansen and Sargent (2022) and Cerreia-Vioglio, Hansen, Maccheroni, and Marinacci (2022).[2]

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

[1] See Diaconis and Freedman (1986) for a classical argument in favor of restricting to a finite-dimensional description of the uncertainty.

[2] It has as a special case the multiplier decision criterion of Hansen and Sargent (2001) axiomatized by Strzalecki (2011).

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 a weighted 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 agent's 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 of a misspecified agent, 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. We say that such an agent is a "statistician type". To support the identification of the statistician 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).[3]

We allow departures from this benchmark with 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 consider the opposite case in which the agent is too lenient in evaluating their model and attributes too much of the unexplained evidence to sampling variability.

We then move to the characterization of the long-run behavior of the 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 terms

[3]Moreover, we observe that SEU maximization, the original robust decision criterion of Hansen and Sargent (2001) and the more recent maxmin robust control of Hansen and Sargent (2022) fail to satisfy these mild rationality requirements.

of relative entropy from 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 statistician type will maintain a non-trivial concern for misspecification. If their behavior converges, it converges to a robust control best reply to the models closer in terms of 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.[4] 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 their 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 there are 20% fewer agents who 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 statistician type is not guaranteed to converge. Indeed, it is possible that their behavior cycles between phases of different misspecification concerns. Still, we characterize the limit action frequency and concern for misspecification. We apply this result to explain the cyclical behavior of monetary policies documented in Sargent (1999) and Sargent (2008). Intuitively, the cycles have the following structure. The agent can play an action whose consequences are well explained by one

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

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 safe action.

We also obtain two axiomatic results that provide a behaviorally 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. At the same time, imposing the standard notion of "more ambiguity averse" of Ghirardato and Marinacci (2002) on the conditional preferences is enough to obtain that for each given period, the preferences are more misspecification concerned the larger the LRT is. More interestingly, the quantitative speed of 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.

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 $u : A \times Y \to \mathbb{R}$ 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.[5]

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 Borel probability measures over outcomes $q^\theta = \left(q_a^\theta\right)_{a \in A} \in \Delta(Y)^A$, 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.

We must impose a few regularity conditions.

Assumption 1. For every $a \in A$:

(i) There exists $\xi_a \in \Delta(Y)$ with $\xi_a \sim p_a^*$ and $\xi_a \sim q_a^\theta$ for all $\theta \in \Theta$, and the densities of p_a^* and q_a^θ with respect to ξ_a are denoted as $\tilde{p}_a^*$ and $\tilde{q}_a^\theta$.[6]

(ii) The maps $\theta \mapsto q_a^\theta$, $(\theta, y) \mapsto \tilde{q}_a^\theta(y)$, and $y \mapsto \tilde{p}_a^*(y)$ are continuous.

(iii) There exists $K > 0$ such that for all $\theta \in \Theta$, $\left|\ln\left(\frac{\tilde{q}_a^\theta}{\tilde{p}_a^*}\right)\right| \leq K$ holds p_a^*-a.s.

Equivalence with a given measure allows us to compute the relevant expectations while allowing for both discrete and continuous outcome spaces. Continuity of the map from parameters to outcome distribution 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 would evaluate an action $a \in A$ according to its (classical) subjective expected utility (Cerreia-Vioglio, Maccheroni, Marinacci, and Montrucchio, 2013b)

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

[5]For every subset C of a Euclidean space, we denote as $\Delta(C)$ the metric space of Borel probability measures on C, endowed with the Prokhorov metric d.

[6]For every $p, q \in \Delta(Y)$, $p \sim q$ means that p and q are mutually absolutely continuous.

[7]Part (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.

That is, they would compute a two-stage expectation of the utility function: they evaluate the utility of the action given the candidate model θ, $\mathbb{E}_{q_a^\theta}\left[u\left(a,y\right)\right]$, 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 terms of relative entropy $R\left(\cdot||\cdot\right)$ to some of the structured models.[8]

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

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

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') 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 distinguishes two levels of uncertainty. At the first level, given a probabilistic model q^θ, the uncertainty about the exact specification of the model is

[8]Recall that for every $p, p' \in \Delta\left(Y\right)$,

$$R\left(p||p'\right) = \begin{cases} \int \log\left(\frac{\mathrm{d}p}{\mathrm{d}p'}\right) \mathrm{d}p & p' \gg p \\ \infty & \text{otherwise} \end{cases}$$

where $p' \gg p$ means that p is absolutely continuous with respect to p'. Appendix A.3 explains under what other distances between probability distributions our convergence results continue to hold.

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

captured by minimizing the expected value of u 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 of uncertainty 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, 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

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

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 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

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

statistic:[11]

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

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 an unambiguous 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 will be invariant to the $N(\Theta)$ choice as long as the following assumption is satisfied.

Assumption 2. (i) $N(\Theta)$ is closed and $p^* \in N(\Theta)$. (ii) There exists $K' > 0$ such that, for all $p \in N(\Theta)$ and $a \in A$, $\xi_a \gg p_a$ and $\ln\left(\frac{\tilde{p}_a}{\tilde{p}^*_a}\right) \leq K'$ holds p^*_a-a.s.

We require two conditions: the unstructured set is sufficiently large to include the actual distribution, and the log of the ratio between the probability density and the actual model can be bounded. This condition 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

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

where $\boldsymbol{\beta} = (\beta_t)_{t \in \mathbb{N}} \in \mathbb{R}^{\mathbb{N}}_{++}$ is an unbounded increasing sequence. A couple of considerations on this formula are in order. First, we do not 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.

[11] The Neyman-Pearson Lemma establishes the performance of the test under correct specification while Foutz and Srivastava (1977) and Vuong (1989) contain the classical results about the informativeness of the likelihood ratio statistics under misspecification. See also Schwartzstein and Sunderam (2021) for a recent economic paper that models agents in a persuasion problem who perform model selection using this statistic.

Second, without further assumptions on $\boldsymbol{\beta}$, how the concern for misspecification reacts to the evidence against or in favor of the structured model is pretty general.[12]

Beyond the likelihood ratio test, a key role is played by the sequence of time scalings $\boldsymbol{\beta}$. Indeed, if the agent is misspecified, the likelihood ratio statistic almost surely diverges to infinity, regardless of the distance between the actual DGP and the set of models in Θ.[13] Therefore, if we want to model agents that adjust their misspecification concern $\lambda^{\boldsymbol{\beta}}(h_t)$ in a way that remains responsive to the model fit, 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.[14] Motivated by these results, we often informally refer to an agent who uses such a time scaling as a "statistician type". Of course, the objective of a statistician can be very different from that of an agent involved in a repeated 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 scaling such that $\beta_t = o(t)$, e.g., $\beta_t = \sqrt{t}$, overly penalizes minor imperfections of the

[12] We discuss in Section 6.3 the maximum level of generality under which our result would continue to hold.

[13] More precisely, the likelihood ratio statistic diverges almost surely if the agent is misspecified about the consequence of some action played infinitely often.

[14] This use of the LRT complements its classical use to decide 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.

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 uncertainty attitudes we expect to observe 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.,[15]

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

Also let

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

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

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

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

A (pure) *policy* is a measurable $\Pi : \mathcal{H} \to A$ that specifies an action for every history. The objective action-contingent probability distribution p^* and a policy Π induce a probability measure $\mathbb{P}_{\Pi}$ on $(A \times Y)^{\mathbb{N}}$.[16] Our interest is not in arbitrary policies, but in policies derived from maximizing the value in equation (1) for some rule $\boldsymbol{\beta}$ determining how the concern for misspecification is adjusted in (2).

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

[15] Throughout the paper, we use the convention $0 \cdot \infty = 0$.

[16] We spell out the $\mathbb{P}_{\Pi}$ derivation in the Appendix.

3.1 Safety and Consistency

We have mentioned that using a linear time scaling $\boldsymbol{\beta}$ has the good statistical property of keeping $\lambda^{\boldsymbol{\beta}}$ 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. Time scaling $\boldsymbol{\beta}$ is ε-*safe* for decision problem (u, A, Y) if for every $\boldsymbol{\beta}$-optimal policy Π and DGP $p^* \in \Delta(Y)^A$

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

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 SEU Bayesian 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. Time scaling $\boldsymbol{\beta}$ is ε-*consistent under almost correct specification* for the exogenous decision problem (u, A, Y) if there exists $\delta > 0$ such that for every $\boldsymbol{\beta}$-optimal policy Π and DGP $p^* \in \Delta(Y)$

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

ε-consistency under almost correct specification requires that sufficiently low levels of misspecification cannot induce considerable ex-post regret.

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 exist an exogenous decision problem and $\varepsilon > 0$ for which there is no ε-safe and ε-consistent under almost correct specification $\boldsymbol{\beta}$ 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 statistician type. In contrast, standard SEU maximization is not safe, while a robust control decision criterion with invariant concern for misspecification a la Hansen and Sargent (2001) or Hansen and Sargent (2022) is not ε-consistent under almost correct specification.[17]

3.2 Long-run Behavior

We are interested in the actions that can arise as the limit behavior of agents with an evolving concern for misspecification. The main results of this section show that we can describe these limit actions through fixed point conditions involving the agent's action, belief, and concern for misspecification. To this end, let $\Theta(a) = \operatorname{argmin}_{\theta\in\Theta} R\left(p_a^* || q_a^\theta\right)$ be the set of distributions 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

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

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

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

3. *Maxmin equilibrium* if

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

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

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

[17] This last observation about 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.

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 in Θ). 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 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 *$\boldsymbol{\beta}$-limit action* if there is a $\boldsymbol{\beta}$-optimal policy Π such that $\mathbb{P}_{\Pi}\left[\sup\{t\colon a_t \neq a\} < \infty\right] > 0$.

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.

Proposition 2. *Suppose that Y is finite and $p^* \in \operatorname{int}\Theta$. Every $\boldsymbol{\beta}$-limit action 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 $\boldsymbol{\beta}$-limit action with no $\theta \in \Theta$ such that $q^{\theta}_{a^*} = p^*_{a^*}$. 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 a worst-case scenario for all the models that are minimally consistent with the structured models.

Finally, if the behavior of the statistician 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 Equilibria Illustrations

The following 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 $z = a + \omega$, where ω is a stochastic term that does not depend on the action and has $\mathbb{E}_{p_a^*}[\omega] = 0$ for all $a \in A$. The agent pays taxes $t = \tau(z) + \varepsilon_1$, where $\tau : \mathbb{R} \to \mathbb{R}$ 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 uncertain true and conjectured aspects of the schedule.*[18] *See Liebman and Zeckhauser (2004) and Rees-Jones and Taubinsky (2020) for the empirical evidence supporting this "schmeduling" bias.*

Simple computations show that $\Theta(a) = \left\{\mathbb{E}_{p_a^*}\left[\frac{\tau(a+\omega)}{a+\omega}\right]\right\}$, i.e., the best fitting marginal taxation is equal to the (lower) average taxation. 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 larger the gap between the average and marginal rate and the higher the excess effort of the agent.

[18] Formally, ε normally distributed implies that Y is not compact, in contrast with the main 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.

In every c-robust equilibrium, this bias is reduced.[19] *To see this observe that since the agent is not perfectly able to explain the equilibrium data, i.e.,* $\min_{\theta\in\Theta} R\left(p_a^*||q_a^\theta\right) > 0$, *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 can prevent, or at least reduce, the counterintuitive prediction that complicated nonlinear taxation codes induce more effort: failures to understand the received tax bill lead to a reduction of 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 rationalizes what is found in 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.*[20]

The following 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* $s \in S$*, 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*

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

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

$$\Theta(a) = \left\{\theta \in \Theta : q_\theta(s) = p^*(s), q_\theta(v) = \frac{p^*(v)}{p^*(\{v : v \leq a\})}, \forall s \in S, \forall v \in V\right\}.$$

[19] See Appendix A.4 for the computations supporting the claim.

[20] In 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.

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. The correlation between valuations and prices results in a positive $\min_{\theta\in\Theta} R\left(p^*_{a^*}||q^\theta_a\right) > 0$ *and makes the agent less confident in their model. Since lower actions are accepted with lower probability, they are less sensitive to the model misspecification, and, therefore, this positive concern leads to even lower bids.*

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

$$\Theta(\alpha) = \text{argmin}_{\theta\in\Theta} \sum_{a\in A} \alpha(a) R\left(p^*_a||q^\theta_a\right)$$

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

$$\text{supp}\nu \subseteq \Theta(\alpha^*), \ \alpha^* \in \Delta\left(BR^\lambda(\nu)\right), \text{ and } \lambda = \min_{\theta\in\Theta} \sum_{a\in A} \alpha^*(a) R\left(p^*_a||q^\theta_a\right)/c.$$

A mixed robust equilibrium allows for mixing 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

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

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

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 $\boldsymbol{\beta}$-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., $\min_{\theta\in\Theta} R\left(p_a^* || q_a^\theta\right)$ is low— but it explains poorly the consequences of a' —i.e., $\min_{\theta\in\Theta} R\left(p_{a'}^* || q_{a'}^\theta\right)$ 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 will oscillate 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 an incentive to break them.[21] 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,

[21] For example, this is what happens under fictitious play.

the agent's preferences converge.[22] 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) and Sargent (2008) and 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 $a = 0$.[23] 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 and normalized to have the same support $[-1, 1]$ and density bounded away from 0. 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 particular, 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 (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

[22] A singleton $\Theta(a)$ is a mild requirement satisfied in many cases, see Fudenberg, Lanzani, and Strack (2021) for a discussion.

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

inflation-specific shocks.[24]

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:

$$\left(\theta_3^* + 1\right)\left(1 - \theta_{1\pi}^*\right) > \theta_0^* + \theta_2^* + \theta_{1a}^*.$$

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 (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 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. We also have some natural comparative statics for the conservative monetary policy frequency in the extreme

[24] An alternative, more explicit specification would have

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

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.

robust equilibria,[25] as a more significant exploitable trade-off between inflation and unemployment induces more time spent using an aggressive monetary policy.

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.[26]

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.[27] 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 $\boldsymbol{\beta}$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

[25] It 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 of misspecification and the conservative policy payoff guarantees equilibrium uniqueness.

[26] In Section 3, we do not want the agent to choose every possible state-contingent function because, for the analysts, what is important is to study the (long-run) behavior of agents in an environment with a given set of actions. Of course, there is no hope of recovering the agent's preferences by looking at their choices between a finite number of alternatives. We need a rich space of alternatives that includes at least all the simple acts.

[27] This 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.

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.

To characterize the agent's dynamic preferences, we consider a collection of binary relations indexed by the observed history. 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 relations are, which is defined using an event E (nonnull and with nonnull complement) and two deterministic and strictly ranked outcomes x and y as measuring rods. Loosely speaking, two relations are (x, y, E) similar if they agree in the evaluation of all pairs of acts that pay x or y on E. With this, an Asymptotic Frequentism axiom singles out the statistician type: for every (x, y, E), the conditional preferences after sufficiently long sequences of outcomes sharing the same empirical frequency must be (x, y, E)-similar. Conversely, a lenient type asymptotically becomes (x, y, E)-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 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 $u : X \to \mathbb{R}$ is a nonconstant affine function, $Q \subseteq \Delta(S)$ is a nonempty set, $\mu \in \Delta(Q)$, $\lambda \geq 0$, and

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

The average robust control representation is the counterpart of (1) when expressed over acts. An apparent difference is the fact that u here takes as input only outcomes instead of pair of actions and consequences. However, this discrepancy is inconsequential, as in Section 3 we can always 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 $\hat{p} \in \Delta\left(\hat{Y}\right)^A$ 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).[28]

5.3 Static Axioms

Our first axiomatic step is a static one. We focus on a particular period, and we characterize in terms of behavioral axioms an agent that evaluates accordingly to equation

[28] See 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 consequence.

(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 $f, g \in \mathcal{F}$, and $x, x' \in X$,*

$$\gamma f + (1-\gamma)\, x \succsim \gamma g + (1-\gamma)\, x \Rightarrow \gamma f + (1-\gamma)\, x' \succsim \gamma g + (1-\gamma)\, x' \qquad \forall \gamma \in (0,1)\,.$$

Continuity. *If $f, g, h \in \mathcal{F}$ the sets*

$$\{\gamma \in [0,1] : \gamma f + (1-\gamma)\, g \succsim h\} \;\; \textit{and} \;\; \{\gamma \in [0,1] : h \succsim \gamma f + (1-\gamma)\, g\}$$

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 $f, g \in \mathcal{F}$ and $\gamma \in (0,1)$*

$$f \sim g \implies \gamma f + (1-\gamma)\, g \succsim f.$$

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

Weak Monotone Continuity. *If $f, g \in \mathcal{F}$, $x \in X$, $(E_i)_{i\in\mathbb{N}} \in \Sigma^{\mathbb{N}}$ with $E_1 \supseteq E_2 \supseteq \ldots$ and $\cap_{n\geq 1} E_n = \emptyset$, then $f \succ g$ implies that there exists $n_0 \in \mathbb{N}$ such that $x E_{n_0} f \succ g$.*

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 facing 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.

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

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

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

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.[29] We say that an event $E \subseteq S$ satisfies the sure-thing principle if, for all $f, g, h, h' \in \mathcal{F}$, $x, y, z, w \in X$ the following conditions are satisfied:

1. If $fEh \succsim gEh$, then $fEh' \succsim gEh'$.
2. If $hEf \succsim hEg$ then $h'Ef \succsim h'Eg$
3. $\frac{x+y}{2} \sim \frac{z+w}{2} \implies \frac{x+y}{2} \sim_E \frac{z+w}{2}$.

We denote by Σ_{st} the set of events that satisfy the sure-thing principle.

Axiom 2 (Structured Savage)**.** *There is a finite set $B \subseteq \Delta(\Omega)$ such that $\Omega \times (\Delta(\Omega) \backslash B)$ is null.* ***P2.*** *$\Sigma_s \subseteq \Sigma_{st}$.* ***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 P2 requires that the agent posits a finite number of models and guarantees that when evaluating acts that only depend on the identity of the structured model, the agent satisfies the Sure-Thing Principle.[30] Structured Savage P4 guarantees

[29] With 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 $f(W \times \Delta(\Omega))g$.

[30] The extension to infinite 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.

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 all $\rho \in \Delta(\Omega)$ and $f, g, h, h' \in \mathcal{F}$,*

$$fWh \succsim_\rho gWh \implies fWh' \succsim_\rho gWh' \qquad \forall W \subseteq \Omega.$$

The Intramodel STP and the Structured Savage's P2 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 $\rho, \rho' \in \Delta(\Omega)$, $f, g \in \mathcal{F}$ and $x \in X$ 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

$$f \succsim_\rho x \iff g \succsim_{\rho'} x.$$

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'.

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

All the agent's structured models have an event with the same probability, are uniform over a model-specific set of outcomes, and at least two are deemed equally likely. It is well-known that equal probability requirements, as in parts (ii) and (iii), are essential for probabilistic sophistication with respect to a 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., to not 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 $x, y \in X$, $\rho, \rho' \in \Delta(\Omega)$ with $x \succ y$, $\Omega \times \{\rho\}$, $\Omega \times \{\rho'\}$ nonnull, and $x_\rho y \sim y_{\rho'} x$. Then $x_\rho y \sim \frac{x+y}{2}{}_{\{\rho,\rho'\}} y$.*

Uncertainty Neutrality over Models guarantees that at the level of bets over models, the agent is "risk-neutral". 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.*

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 averse 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 $\left(\succsim^h\right)_{h\in\mathcal{H}}$ indexed by the realized history. Recall that the relevant set of length t 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)**.** *For every $h \in \mathcal{H}$, every $\omega \in \Omega$, and for every $f, g \in \mathcal{F}$ that are Σ_s-measurable*

$$f \succsim^h_{\{\omega\}} g \Longleftrightarrow f \succsim^{(h,\omega)} g.$$

The second dynamic axiom requires Bayesian rationality when considering acts whose consequence only depends on the structured model. To link the misspecification concern to the evidence, we need to translate the notion of the likelihood ratio test for a set of models in our setting:

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

Axiom 8 (Q-Likelihood)**.** *There exists $N\left(\left\{\rho_q : q \in Q\right\}\right) \subseteq \Delta\left(\Omega\right)$ such that for every $t \in \mathbb{N}$, $h_t, h'_t \in \mathcal{H}_t$, if*

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

then for every $q \in Q$, $\succsim^{h_t}_{\rho_q}$ is more concerned with misspecification than $\succsim^{h'_t}_{\rho_q}$.

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 $\succsim_{\rho_q}$, as different posteriors over structured models imply that the unconditional preferences are generally not rankable in terms of more misspecification aversion.

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

1. *Each $\succsim^h$ satisfies the axioms of Theorem 3,*

2. *For each* $\succsim^h$,

$$\left\{\rho_q : q \in Q\right\} \subseteq \operatorname{int} \Delta(\Omega)$$

is the smallest closed set $B \subseteq \Delta(\Omega)$ *such that* $\Omega \times (\Delta(\Omega) \setminus B)$ *is* $\succsim^h$*-null,*

3. $\left(\succsim^h\right)_{h \in \mathcal{H}}$ *satisfies Constant Preference Invariance, Q-Likelihood, and Dynamic Consistency over Models.*

Then each $\succsim^h$ *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.*

If the agent observes a probability 0 event, Bayes rule has no bite, and the agent can (and should) revise their beliefs more radically than the adjustment in our model. Therefore, Proposition 4 and the following results focus on the case where each ρ_q has full support. That is the same case considered in the learning part of the paper.[31]

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.

For every $x, y \in X$ and $E \in \Sigma$ the preference relation $\succsim$ is (x, y, E)*-similar* to $\succsim'$ if

$$fEx \succ' gEy \implies fEx \succ gEy \qquad \forall f, g \in \mathcal{F}$$

and

$$z \succ' w \iff z \succ w \qquad \forall z, w \in X.$$

Axiom 9 (Asymptotic Frequentism)**.** *For every* $\rho \in \Delta(\Omega)$ *there exists* $\succsim^{\lim}$ *with an average robust control representation such that for all* $x, y \in X$ *and* $E \in \Sigma$ *with* $x \succ^{\emptyset} y$ *and* $S \setminus E$ *nonnull the following holds. There is* $t \in \mathbb{N}$ *such that if* $t \geq \tau$ *and* h_t *has outcome frequency* ρ *then* $\succsim^{h_t}$ *is* (x, y, E)*-similar to* $\succsim^{\lim}$.

The axiom requires that for every measuring triplet (x, y, E), sufficiently long sequences of outcomes with the same empirical frequency stabilize preferences.

[31] To be more precise, in the learning model, the data-generating processes conditional on every action have the same support $\tilde{Y} \subseteq Y$, and that support has objective probability 1. Since both the subjective and the objective models are concentrated on the same set, this is equivalent to full support on the restricted consequence space $\tilde{Y}$.

Proposition 5. *If $\left(\succsim^h\right)_{h\in\mathcal{H}}$ 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 that is not in $\{\rho_q : q \in Q\}$,*

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

exists and is finite.

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 the requirement 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 statistician type.

The proof of the result has three 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) is converging. 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 to a finite value. 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 and an outcome z strictly worse than y such that the preference with the high concern for misspecification limit strictly prefers y for sure to xEz. The opposite holds for the preference with the low concern limit. The final step rules out divergence to plus infinity by showing that, in that case, the preferences would converge to a maxmin model on the unrestricted set of distributions and that such preferences cannot be (x, y, E) similar to any finite level of misspecification 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 and $\Omega \times \{\rho\}$ is $\succsim^{\emptyset}$-nonnull. 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 $\hat{\rho}$, $x \succsim^{h_t} f$.*

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 $\left(\succsim^h\right)_{h\in\mathcal{H}}$ 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*

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

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* $x, y \in X$*,* $E \in \Sigma$*,* $\rho \in \Delta(\Omega)$*,* $B \subseteq \Delta(\Omega)$*, be such that* $x \succ y$ *and* $S \setminus E$*,* $\Omega \times B$ *are nonnull. For every SEU preference* $\succeq$ *such that* $\succeq_B$ *is less misspecification averse than* $\succsim_B$*, there exists* $\tau \in \mathbb{N}$ *such that for all* $t \geq \tau$ *and all* h_t *with outcome frequency* ρ*,* $\succsim^{h_t}$ *is* (x, y, E)*-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)-similar to the one updated preferences of an SEU whose model contingent preferences were initially less misspecification averse than the agent.

Proposition 7. *If* $\left(\succsim^h\right)_{h \in \mathcal{H}}$ *satisfies the assumptions of Proposition 4 and Asymptotic Leniency then for every sequence* $\left(h_{t_n}\right)_{n \in \mathbb{N}}$ *with constant outcome frequency*

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

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

"lightbulb realizations" or evolutionary pressure. The key difference with our approach is that in these papers, 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. 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.

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 Fudenberg, Lanzani, and Strack (2021), Bohren and Hauser (2021), Esponda, Pouzo, and Yamamoto (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 Number follows the formalization of this bias proposed by Rabin (2002).

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.

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

[32]Given 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).

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 full revealed-preference elicitation of the relevant probability distributions, viewed as subjective statistical models. See also Dean and Ortoleva (2017) for a less related decision where the agent has a prior over multiple data-generating processes and evaluates each of them with rank-dependent utility.

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 Number with uncertainty aversion. However, the findings in Esponda, Vespa, and Yuksel (2020) 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 drifts towards the safest action as repeated observations about the shortcomings of their model are accumulated. Notably, the safe action is not the best reply to observed empirical frequency, which suggests that, as in our model, even in the medium run, the agents do not altogether drop their models; they just 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.

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 $F : \mathbb{N} \times \mathbb{R}_+ \to \mathbb{R}_+$ is a function that maps a period and a realization of the likelihood-ratio statistic into a level of misspecification concern. The primary learning results extend, with the statistician behavior pinned down by the condition that for each $k > 0$,

$\lim_{n\to\infty} F(n, kn)$ exists and is finite, 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., for 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 changes, little is changed. However, analyzing an uncommitted, forward-looking, and sophisticated agent playing an intra-personal equilibrium with their future selves would require novel tools and results.

6.5 Endogenous Structured Models

The more natural extension for the axiomatic 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 conditions 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.

Definition 12. A nonempty set $Q \subseteq \Delta(S)$ is *identifiable* if there exists a measurable function $k : S \to Q$ such that

$$q(\{s : k(s) = q\}) = 1 \qquad \forall q \in Q.$$

In that case, we say that k identifies Q.

When considering a structured environment, we required identifiability through the map $(\omega, \rho) \mapsto \rho \times \delta_\rho$. However, in addition, 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 level of 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 novel 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 $P = \Delta(Y)^A$ be the space of all action-dependent outcome distributions,

$$p_a \in \Delta(Y)$$

denote the a-th component of $p \in P$, and

$$||p - p'|| = \max_{a \in A} d\left(p_a, p'_a\right) \qquad \forall p, p' \in P.$$

For an arbitrary Borel measurable subset K of a Euclidean space, we endow the space $K^{\mathbb{N}}$ with the Borel σ-algebra, $\mathcal{B}\left(K^{\mathbb{N}}\right)$, corresponding to the product topology on $K^{\mathbb{N}}$. We denote by $k^t = (k_1, ..., k_t) \in K^t$ both the finite sequence and the elementary cylinder in $K^{\mathbb{N}}$ that it identifies. We denote by $k^\infty = (k_1, ..., k_t, ...)$ an arbitrary element of $K^{\mathbb{N}}$.

For every policy $\Pi \in A^{\mathcal{H}}$, the objective probability distribution over infinite histories is defined over elementary cylinders as

$$\mathbb{P}_\Pi\left(a^t, y^t\right) = \begin{cases} \prod_{\tau=1}^{t} p^*_{a_\tau}\left(y_\tau\right) & \text{if for all } \tau \in \{1, ..., t\} \text{ we have } a_{\tau+1} = \Pi\left(a^\tau, y^\tau\right) \\ 0 & \text{otherwise.} \end{cases} \tag{6}$$

Since this set of finite-dimensional distributions is consistent, there is a unique probability distribution over infinite sequences of action-outcome pairs with these marginals,

defined through the Kolmogorov Extension Theorem.

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

$$p_a^{h_t}(y) = \frac{\sum_{\tau=1}^{t} \mathbb{I}_{(a,y)}(a_\tau, y_\tau)}{\sum_{\tau=1}^{t} \mathbb{I}_a(a_\tau)} \qquad \forall y \in Y$$

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

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

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

Proof of Lemma 1. Since u is continuous, $p_a \mapsto \mathbb{E}_{p_a}[u(a,y)]$ is continuous, and $p_a \mapsto R\left(p_a || q_a^\theta\right)$ is lower semicontinuous by Theorem 1.47 in Liese and Vajda (1987). Therefore, the set $E := \left\{p_a \in \Delta(Y) : \mathbb{E}_{p_a}[u(a,y)] + \frac{1}{\lambda} R\left(p_a || q_a^\theta\right) \leq \mathbb{E}_{q_a^\theta}[u(a,y)]\right\}$ is closed, and we clearly have

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

Since Y is compact, by Theorem 15.11 in Aliprantis and Border (2013) so are $\Delta(Y)$ and $\Delta(Y) \cap E$. Since $p_a \mapsto \mathbb{E}_{p_a}[u(a,y)] + \frac{1}{\lambda} R\left(p_a || q_a^\theta\right)$ is real-valued and lower semicontinuous on the compact $\Delta(Y) \cap E$, it admits a minimizer by Theorem 2.43 in Aliprantis and Border (2013). ■

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*

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

Proof of Lemma 2. We have

$$\begin{aligned}
& \lambda^{\boldsymbol{\beta}}(h_t) = -\log\left(\frac{\max_{\theta \in \Theta} \prod_{\tau=1}^{t} \tilde{q}_{a_\tau}^{\theta}(y_\tau)}{\max_{r \in N(\Theta)} \prod_{\tau=1}^{t} \tilde{r}_{a_\tau}(y_\tau)}\right) / \beta_t \\
= \; & \log\left(\frac{\max_{r \in N(\Theta)} \prod_{\tau=1}^{t} \tilde{r}_{a_\tau}(y_\tau)}{\max_{\theta \in \Theta} \prod_{\tau=1}^{t} \tilde{q}_{a_\tau}^{\theta}(y_\tau)}\right) / \beta_t = \left(\log\left(\frac{\prod_{\tau=1}^{t} \tilde{p}_{a_\tau}(y_\tau)}{\prod_{\tau=1}^{t} \tilde{q}_{a_\tau}^{\hat{\theta}}(y_\tau)}\right)\right) / \beta_t \\
= \; & \left(\log\left(\prod_{\tau=1}^{t} \tilde{p}_{a_\tau}(y_\tau)\right) - \log\left(\prod_{\tau=1}^{t} \tilde{q}_{a_\tau}^{\hat{\theta}}(y_\tau)\right)\right) / \beta_t \\
= \; & \frac{\log\left(\prod_{y \in Y} \prod_{a \in A} \tilde{p}_a(y)^{\sum_{\tau=1}^{t} \mathbb{I}_a(a_\tau) p_a^{h_t}(y)}\right) - \log\left(\prod_{y \in Y} \prod_{a \in A} \tilde{q}_a^{\hat{\theta}}(y)^{\sum_{\tau=1}^{t} \mathbb{I}_a(a_\tau) p_a^{h_t}(y)}\right)}{\beta_t} \\
= \; & \frac{\begin{array}{c} \sum_{a \in A} \sum_{\tau=1}^{t} \mathbb{I}_a(a_\tau) \sum_{y \in Y} p_a^{h_t}(y) \log(\tilde{p}_a(y)) \\ - \sum_{a \in A} \sum_{\tau=1}^{t} \mathbb{I}_a(a_\tau) \sum_{y \in Y} p_a^{h_t}(y) \log\left(\tilde{q}_a^{\hat{\theta}}(y)\right) \end{array}}{\beta_t} \\
= \; & \frac{\sum_{a \in A} \sum_{\tau=1}^{t} \mathbb{I}_a(a_\tau) \left(\sum_{y \in Y} p_a^{h_t}(y) \log(\tilde{p}_a(y)) - \sum_{y \in Y} p_a^{h_t}(y) \log\left(\tilde{q}_a^{\hat{\theta}}(y)\right)\right)}{\beta_t} \\
= \; & \sum_{a \in A} \sum_{\tau=1}^{t} \mathbb{I}_a(a_\tau) \int_Y \log\left(\frac{\tilde{p}_a(y)}{\tilde{q}_a^{\hat{\theta}}(y)}\right) \mathrm{d}p_a^{h_t}(y) / \beta_t.
\end{aligned}$$

■

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 all* $a \in A$ *and* $\theta \in \Theta$ *we have*

$$\lim_{\lambda \downarrow 0} \min_{p_a \in \Delta(Y)} \mathbb{E}_{p_a}[u(a,y)] + \frac{1}{\lambda} R\left(p_a || q_a^{\theta}\right) = \mathbb{E}_{q_a^{\theta}}[u(a,y)].$$

Proof. Since $\mathbb{E}_{p_a}[u(a,y)] + \frac{1}{\lambda} R\left(p_a || q_a^{\theta}\right)$ is weakly decreasing in λ for all $p_a \in \Delta(Y)$,

also $\min_{p_a \in \Delta(Y)} \mathbb{E}_{p_a}\left[u\left(a,y\right)\right] + \frac{1}{\lambda} R\left(p_a || q_a^\theta\right)$ is weakly decreasing in λ. Moreover, we have

$$\limsup_{\lambda \downarrow 0} \min_{p_a \in \Delta(Y)} \mathbb{E}_{p_a}\left[u\left(a,y\right)\right] + \frac{1}{\lambda} R\left(p_a || q_a^\theta\right) \leq \limsup_{\lambda \downarrow 0} \mathbb{E}_{q_a^\theta}\left[u\left(a,y\right)\right] + \frac{1}{\lambda} 0 = \mathbb{E}_{q_a^\theta}\left[u\left(a,y\right)\right]$$

and the one considered in the statement is a weakly increasing sequence bounded from above. Therefore, the limit exists and the statement is proved by contradiction, showing that we cannot have

$$\lim_{\lambda \downarrow 0} \min_{p_a \in \Delta(Y)} \mathbb{E}_{p_a}\left[u\left(a,y\right)\right] + \frac{1}{\lambda} R\left(p_a || q_a^\theta\right) < \mathbb{E}_{q_a^\theta}\left[u\left(a,y\right)\right].$$

If that was the case, there would be a sequence $\left(\lambda_n\right)_{n \in \mathbb{N}} \in \mathbb{R}_{++}^\infty$ and an $\varepsilon \in \mathbb{R}_{++}$ with

$$\lim_{n \to \infty} \lambda_n = 0 \text{ and } \lim_{n \to \infty} \min_{p_a \in \Delta(Y)} \mathbb{E}_{p_a}\left[u\left(a,y\right)\right] + \frac{1}{\lambda_n} R\left(p_a || q_a^\theta\right) = \mathbb{E}_{q_a^\theta}\left[u\left(a,y\right)\right] - \varepsilon. \tag{7}$$

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

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

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

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

Therefore, there exist $M \in \mathbb{R}$ and $K \in \mathbb{R}_{++}$ such that

$$\begin{aligned} & \hat{p}_a\left(\left\{y \in Y : u\left(a,y\right) \leq M - K\right\}\right) - q_a^\theta\left(\left\{y \in Y : u\left(a,y\right) \leq M\right\}\right) \\ \geq & \ \varepsilon/2 \left[\max_{y \in Y} u\left(a,y\right) - \min_{y \in Y} u\left(a,y\right)\right]. \end{aligned}$$

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

$$\begin{aligned} & p_a^n\left(\{y \in Y : u(a,y) \leq M\}\right) - q_a^\theta\left(\{y \in Y : u(a,y) \leq M\}\right) \\ \geq \; & \varepsilon/4 \left[\max_{y \in Y} u(a,y) - \min_{y \in Y} u(a,y)\right]. \end{aligned}$$

But then

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

where the inequality follows by Theorem 1.30 in Liese and Vajda (1987). But the last term diverges to $+\infty$ as n goes to infinity, a contradiction with

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

■

Lemma 4. *(i) For all $a \in A$, the function $G : \Delta(\Theta) \times \mathbb{R}_+ \to \mathbb{R}$ defined by*

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

and

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

is continuous.

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

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

is upper hemicontinuous.

Proof. (i) Fix $a \in A$. For every $\theta \in \Theta$, let $F(\theta, 0) := \mathbb{E}_{q_a^\theta}\left[u(a,y)\right]$ and observe that

for each $\lambda \in \mathbb{R}_{++}$, by Proposition 1.4.2 in Dupuis and Ellis (2011) we have

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

Since Y is compact and u is continuous the RHS is bounded by

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

Therefore, 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 so that

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

and

$$F(\cdot, \lambda) = -\frac{1}{\lambda} \log\left(\int_Y \exp(-\lambda u(a, y)) \, \mathrm{d}q_a^{(\cdot)}(y)\right)$$

are continuous by Assumption 1 (ii) and the definition of weak convergence of measures. Moreover, $F(\theta, \cdot)$ is continuous by Lemma 3. We now show that for every $K \in \mathbb{R}_{++}$ the class $(F(\cdot, \lambda) : \Theta \to \mathbb{R})_{\lambda \in [0, K]}$ is equicontinuous and therefore F is jointly continuous by Lemma 6 in Dunford and Schwartz (1988). Let $\hat{\theta} \in \Theta$, $\lambda \in (0, K]$, and $\varepsilon \in \mathbb{R}_{++}$. Since Y is compact and u is continuous it is uniformly continuous and there exists $\varepsilon' \in \mathbb{R}_{++}$ such that for all $y, y' \in Y$,

$$||y - y'||_\infty \leq \varepsilon' \implies |u(a, y) - u(a, y')| \leq \frac{\varepsilon}{2}.$$

Moreover, since

$$\lim_{\lambda \to 0} \frac{(\exp(-\lambda \min_{y \in Y} u(a, y)) - \exp(-\lambda \max_{y \in Y} u(a, y)))}{\lambda} = \max_{y \in Y} u(a, y) - \min_{y \in Y} u(a, y)$$

ε' can be chosen to be so small that

$$\frac{\varepsilon' (\exp(-\lambda \min_{y \in Y} u(a, y)) - \exp(-\lambda \max_{y \in Y} u(a, y)))}{\exp(-\lambda \min_{y \in Y} u(a, y))} \leq \lambda \frac{\varepsilon}{2} \qquad \lambda \in (0, K]. \tag{8}$$

Similarly, since Θ is compact and $\theta \mapsto q_a^\theta$ is continuous by Assumption 1 (ii), it is also

uniformly continuous. Therefore, there is a $\delta \in \mathbb{R}_{++}$ such that

$$\left|\left|\theta - \hat{\theta}\right|\right|_{\infty} \leq \delta \implies d\left(q_a^{\theta}, q_a^{\hat{\theta}}\right) \leq \varepsilon'. \tag{9}$$

Suppose without loss of generality that

$$\int_Y \exp\left(-\lambda u\left(a,y\right)\right) \mathrm{d}q_a^{\theta}\left(y\right) \geq \int_Y \exp\left(-\lambda u\left(a,y\right)\right) \mathrm{d}q_a^{\hat{\theta}}\left(y\right).$$

The right hand side of (9) implies that

$$\left| \begin{array}{c} -\frac{1}{\lambda}\log\left(\int_Y \exp\left(-\lambda u\left(a,y\right)\right) \mathrm{d}q_a^{\theta}\left(y\right)\right) \\ +\frac{1}{\lambda}\log\left(\int_Y \exp\left(-\lambda u\left(a,y\right)\right) \mathrm{d}q_a^{\hat{\theta}}\left(y\right)\right) \end{array} \right| = \frac{1}{\lambda}\log\left(\frac{\int_Y \exp\left(-\lambda u\left(a,y\right)\right) \mathrm{d}q_a^{\theta}\left(y\right)}{\int_Y \exp\left(-\lambda u\left(a,y\right)\right) \mathrm{d}q_a^{\hat{\theta}}\left(y\right)}\right)$$

$$= \frac{1}{\lambda}\log\left(\frac{\int_{\exp\left(-\lambda \max_{y\in Y} u(a,y)\right)}^{\exp\left(-\lambda \min_{y\in Y} u(a,y)\right)} \left(1 - q_a^{\theta}\left(\left\{y : \exp\left(-\lambda u\left(a,y\right)\right) \leq r\right\}\right)\right) \mathrm{d}r + \exp\left(-\lambda \max_{y\in Y} u\left(a,y\right)\right)}{\int_{\exp\left(-\lambda \max_{y\in Y} u(a,y)\right)}^{\exp\left(-\lambda \min_{y\in Y} u(a,y)\right)} \left(1 - q_a^{\hat{\theta}}\left(\left\{y : \exp\left(-\lambda u\left(a,y\right)\right) \leq r\right\}\right)\right) \mathrm{d}r + \exp\left(-\lambda \max_{y\in Y} u\left(a,y\right)\right)}\right)$$

$$\leq \frac{1}{\lambda}\log\left(\frac{\int_{\exp\left(-\lambda \max_{y\in Y} u(a,y)\right)}^{\exp\left(-\lambda \min_{y\in Y} u(a,y)\right)} \left(1 - q_a^{\theta}\left(\left\{y : \exp\left(-\lambda u\left(a,y\right)\right) \leq r\right\}\right)\right) \mathrm{d}r}{\int_{\exp\left(-\lambda \max_{y\in Y} u(a,y)\right)}^{\exp\left(-\lambda \min_{y\in Y} u(a,y)\right)} \left(1 - q_a^{\hat{\theta}}\left(\left\{y : \exp\left(-\lambda u\left(a,y\right)\right) \leq r\right\}\right)\right) \mathrm{d}r}\right)$$

$$= \frac{1}{\lambda}\log\left(\frac{\int_{\exp\left(-\lambda \max_{y\in Y} u(a,y)\right)}^{\exp\left(-\lambda \min_{y\in Y} u(a,y)\right)} \left(1 - q_a^{\theta}\left(\left\{y : u\left(a,y\right) \geq -\frac{\log r}{\lambda}\right\}\right)\right) \mathrm{d}r}{\int_{\exp\left(-\lambda \max_{y\in Y} u(a,y)\right)}^{\exp\left(-\lambda \min_{y\in Y} u(a,y)\right)} \left(1 - q_a^{\hat{\theta}}\left(\left\{y : \exp\left(-\lambda u\left(a,y\right)\right) \leq r\right\}\right)\right) \mathrm{d}r}\right)$$

$$\leq \frac{1}{\lambda}\log\left(\frac{\int_{\exp\left(-\lambda \max_{y\in Y} u(a,y)\right)}^{\exp\left(-\lambda \min_{y\in Y} u(a,y)\right)} \left(1 - q_a^{\hat{\theta}}\left(\left\{y : u\left(a,y\right) - \frac{\varepsilon}{2} \geq -\frac{\log r}{\lambda}\right\}\right) + \varepsilon'\right) \mathrm{d}r}{\int_{\exp\left(-\lambda \max_{y\in Y} u(a,y)\right)}^{\exp\left(-\lambda \min_{y\in Y} u(a,y)\right)} \left(1 - q_a^{\hat{\theta}}\left(\left\{y : \exp\left(-\lambda u\left(a,y\right)\right) \leq r\right\}\right)\right) \mathrm{d}r}\right)$$

$$= \frac{1}{\lambda}\log\left(\frac{\begin{array}{c}\int_{\exp\left(-\lambda \max_{y\in Y} u(a,y)\right)}^{\exp\left(-\lambda \min_{y\in Y} u(a,y)\right)} \left(1 - q_a^{\hat{\theta}}\left(\left\{y : \exp\left(-\lambda u\left(a,y\right)\right) \leq r\exp\left(-\lambda\frac{\varepsilon}{2}\right)\right\}\right)\right) \mathrm{d}r \\ +\varepsilon'\left(\exp\left(-\lambda \min_{y\in Y} u\left(a,y\right)\right) - \exp\left(-\lambda \max_{y\in Y} u\left(a,y\right)\right)\right)\end{array}}{\int_{\exp\left(-\lambda \max_{y\in Y} u(a,y)\right)}^{\exp\left(-\lambda \min_{y\in Y} u(a,y)\right)} \left(1 - q_a^{\hat{\theta}}\left(\left\{y : \exp\left(-\lambda u\left(a,y\right)\right) \leq r\right\}\right)\right) \mathrm{d}r}\right)$$

$$
= \frac{1}{\lambda} \log \left(\frac{\begin{array}{c} \int_{\exp\left(-\lambda\left(\max_{y\in Y} u(a,y)-\frac{\varepsilon}{2}\right)\right)}^{\exp\left(-\lambda\left(\min_{y\in Y} u(a,y)-\frac{\varepsilon}{2}\right)\right)} \left(1 - q_a^{\hat{\theta}}\left(\{y : \exp\left(-\lambda u\left(a,y\right)\right) \leq r\}\right)\right) \exp\left(\lambda \frac{\varepsilon}{2}\right) \mathrm{d}r \\ +\varepsilon' \left(\exp\left(-\lambda \min_{y\in Y} u\left(a,y\right)\right) - \exp\left(-\lambda \max_{y\in Y} u\left(a,y\right)\right)\right) \end{array}}{\int_{\exp\left(-\lambda \max_{y\in Y} u(a,y)\right)}^{\exp\left(-\lambda \min_{y\in Y} u(a,y)\right)} \left(1 - q_a^{\hat{\theta}}\left(\left\{y : u\left(a,y\right) \geq -\frac{\log r}{\lambda}\right\}\right)\right) \mathrm{d}r} \right)
$$

$$
= \frac{1}{\lambda} \log \left(\frac{\begin{array}{c} \exp\left(\lambda \frac{\varepsilon}{2}\right) \int_{\exp\left(-\lambda\left(\max_{y\in Y} u(a,y)-\frac{\varepsilon}{2}\right)\right)}^{\exp\left(-\lambda\left(\min_{y\in Y} u(a,y)-\frac{\varepsilon}{2}\right)\right)} \left(1 - q_a^{\hat{\theta}}\left(\{y : \exp\left(-\lambda u\left(a,y\right)\right) \leq r\}\right)\right) \mathrm{d}r \\ +\varepsilon' \left(\exp\left(-\lambda \min_{y\in Y} u\left(a,y\right)\right) - \exp\left(-\lambda \max_{y\in Y} u\left(a,y\right)\right)\right) \end{array}}{\int_{\exp\left(-\lambda \max_{y\in Y} u(a,y)\right)}^{\exp\left(-\lambda \min_{y\in Y} u(a,y)\right)} \left(1 - q_a^{\hat{\theta}}\left(\left\{y : u\left(a,y\right) \geq -\frac{\log r}{\lambda}\right\}\right)\right) \mathrm{d}r} \right)
$$

$$
= \frac{1}{\lambda} \log \left(\frac{\begin{array}{c} \exp\left(\lambda \frac{\varepsilon}{2}\right) \int_{\exp\left(-\lambda\left(\max_{y\in Y} u(a,y)-\frac{\varepsilon}{2}\right)\right)}^{\exp\left(-\lambda \min_{y\in Y} u(a,y)\right)} \left(1 - q_a^{\hat{\theta}}\left(\{y : \exp\left(-\lambda u\left(a,y\right)\right) \leq r\}\right)\right) \mathrm{d}r \\ +\varepsilon' \left(\exp\left(-\lambda \min_{y\in Y} u\left(a,y\right)\right) - \exp\left(-\lambda \max_{y\in Y} u\left(a,y\right)\right)\right) \end{array}}{\int_{\exp\left(-\lambda \max_{y\in Y} u(a,y)\right)}^{\exp\left(-\lambda \min_{y\in Y} u(a,y)\right)} \left(1 - q_a^{\hat{\theta}}\left(\left\{y : u\left(a,y\right) \geq -\frac{\log r}{\lambda}\right\}\right)\right) \mathrm{d}r} \right)
$$

$$
= \frac{1}{\lambda} \log \left(\frac{\begin{array}{c} \exp\left(\lambda \frac{\varepsilon}{2}\right) \left(\begin{array}{c} \int_{\exp\left(-\lambda\left(\max_{y\in Y} u(a,y)\right)\right)}^{\exp\left(-\lambda\left(\min_{y\in Y} u(a,y)\right)\right)} \left(1 - q_a^{\hat{\theta}}\left(\{y : \exp\left(-\lambda u\left(a,y\right)\right) \leq r\}\right)\right) \mathrm{d}r \\ - \int_{\exp\left(-\lambda\left(\max_{y\in Y} u(a,y)\right)\right)}^{\exp\left(-\lambda\left(\max_{y\in Y} u(a,y)-\frac{\varepsilon}{2}\right)\right)} \left(1 - q_a^{\hat{\theta}}\left(\{y : \exp\left(-\lambda u\left(a,y\right)\right) \leq r\}\right)\right) \end{array} \right) \\ +\varepsilon' \left(\exp\left(-\lambda \min_{y\in Y} u\left(a,y\right)\right) - \exp\left(-\lambda \max_{y\in Y} u\left(a,y\right)\right)\right) \end{array}}{\int_{\exp\left(-\lambda \max_{y\in Y} u(a,y)\right)}^{\exp\left(-\lambda \min_{y\in Y} u(a,y)\right)} \left(1 - q_a^{\hat{\theta}}\left(\left\{y : u\left(a,y\right) \geq -\frac{\log r}{\lambda}\right\}\right)\right) \mathrm{d}r} \right)
$$

$$
\leq \frac{1}{\lambda} \log \left(\frac{\begin{array}{c} \exp\left(\lambda \frac{\varepsilon}{2}\right) \left(\int_{\exp\left(-\lambda\left(\max_{y\in Y} u(a,y)\right)\right)}^{\exp\left(-\lambda\left(\min_{y\in Y} u(a,y)\right)\right)} \left(1 - q_a^{\hat{\theta}}\left(\{y : \exp\left(-\lambda u\left(a,y\right)\right) \leq r\}\right)\right) \mathrm{d}r \right) \\ +\varepsilon' \left(\exp\left(-\lambda \min_{y\in Y} u\left(a,y\right)\right) - \exp\left(-\lambda \max_{y\in Y} u\left(a,y\right)\right)\right) \end{array}}{\int_{\exp\left(-\lambda \max_{y\in Y} u(a,y)\right)}^{\exp\left(-\lambda \min_{y\in Y} u(a,y)\right)} \left(1 - q_a^{\hat{\theta}}\left(\left\{y : u\left(a,y\right) \geq -\frac{\log r}{\lambda}\right\}\right)\right) \mathrm{d}r} \right)
$$

$$
\begin{aligned}
&= \frac{1}{\lambda}\log\left(\exp\left(\lambda\frac{\varepsilon}{2}\right)+\frac{\varepsilon'\left(\exp\left(-\lambda\min_{y\in Y}u\left(a,y\right)\right)-\exp\left(-\lambda\max_{y\in Y}u\left(a,y\right)\right)\right)}{\int_{\exp\left(-\lambda\max_{y\in Y}u(a,y)\right)}^{\exp\left(-\lambda\min_{y\in Y}u(a,y)\right)}\left(1-q_a^{\hat{\theta}}\left(\left\{y:u\left(a,y\right)\geq-\frac{\log r}{\lambda}\right\}\right)\right)\mathrm{d}r}\right)\\
&\leq \frac{1}{\lambda}\log\left(\exp\left(\lambda\frac{\varepsilon}{2}\right)+\frac{\varepsilon'\left(\exp\left(-\lambda\min_{y\in Y}u\left(a,y\right)\right)-\exp\left(-\lambda\max_{y\in Y}u\left(a,y\right)\right)\right)}{\exp\left(-\lambda\min_{y\in Y}u\left(a,y\right)\right)}\right)\\
&\leq \frac{1}{\lambda}\log\left(\exp\left(\lambda\frac{\varepsilon}{2}\right)+\lambda\frac{\varepsilon}{2}\right)
\end{aligned}
$$

where the second inequality follows from the definition of the Prokhorov metric, the fifth equality by the change of variable theorem (see Theorem 13.46 in Aliprantis and Border (2013)), and the fourth equality follows by equation (8). However, the last term is bounded for $\lambda\in(0,K]$, and so we have proved the desired equicontinuity. Since we have shown equicontinuity of the elements of the family except $F\left(\cdot,0\right)$ and the latter is continuous, we have proved the equicontinuity of the entire family.

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

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

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

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

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

(ii) Follows by (i) and Theorem 17.31 in Aliprantis and Border (2013). ■

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 a α^* weighted average of the relative entropy from the true DGP.

Lemma 5. *Let $\left(a_t,y_t\right)_{t\in\mathbb{N}}\in\left(A\times Y\right)^{\mathbb{N}}$ be such that $\tilde{p}^*_{a_t}\left(y_t\right)>0$ for all $t\in\mathbb{N}$ and let $\alpha^*\in\Delta(A)$. For each $t\in\mathbb{N}$, set $h_t=\left(a^t,y^t\right)$, and let $\hat{\theta}\left(h_t\right)$ and $r\left(h_t\right)$ be two arbitrary*

selections from $\operatorname{argmax}_{\theta\in\Theta}\prod_{\tau=1}^{t}\tilde{p}_{a_\tau}^{\theta}(y_\tau)$ *and* $\operatorname{argmax}_{p\in N(\Theta)}\prod_{\tau=1}^{t}\tilde{p}_{a_\tau}(y_\tau)$ *respectively. If*

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

then

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

Moreover,

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

Proof. Fix $a\in A$ and consider a sequence $\left(\theta_n,p_a^n\right)_{n\in\mathbb{N}}\in\left(\Theta\times\Delta(Y)\right)^{\mathbb{N}}$ with

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

and $\operatorname{supp}p_a^n\subseteq\{y:\tilde{p}_a^*(y)>0\}$. By Assumption 1 (ii-iii) the assumptions of Theorem 15.7.3 in Kallenberg (1973) are satisfied for the sequence of integrand functions and probability measures $\left(\log\left(\frac{\tilde{q}_a^{\theta_n}(\cdot)}{\tilde{p}_a^*(\cdot)}\right),p_a^n\right)_{n\in\mathbb{N}}$. Therefore, we have

$$\lim_{n\to\infty}-\int_Y\log\left(\frac{\tilde{q}_a^{\theta^n}(y)}{\tilde{p}_a^*(y)}\right)\mathrm{d}p_a^n(y)=-\int_Y\log\left(\frac{\tilde{q}_a^{\bar{\theta}}(y)}{\tilde{p}_a^*(y)}\right)\mathrm{d}\bar{p}_a(y).$$

Since Θ and $\Delta(Y)^A$ are metric spaces, by linearity in the finite dimensional α this implies that

$$(p,\theta,\alpha)\mapsto-\sum_{a\in A}\alpha(a)\int_Y\log\left(\frac{\tilde{q}_a^{\theta}(y)}{\tilde{p}_a^*(y)}\right)\mathrm{d}p_a(y)\qquad p\in\Delta(Y)^A \tag{10}$$

is continuous on

$$\left\{p\in\Delta(Y)^A:\forall a\in A,\operatorname{supp}p_a\subseteq\{y:\tilde{p}_a^*(y)>0\}\right\}\times\Theta\times\Delta(A).$$

Observe that, as by assumption for each $t \in \mathbb{N}$, $p^*_{a_t}(y_t) > 0$, we have

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

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

$$\int_Y -\log\left(\frac{\tilde{r}_a(h_t)(y)}{\tilde{p}_a^*(y)}\right)\mathrm{d}p_a^{h_t}(y) \leq 0 \qquad \forall t \in \mathbb{N}.$$

Also,

$$0 \leq \int_Y -\log\left(\frac{\tilde{r}_a(h_t)(y)}{\tilde{p}_a^*(y)}\right)\mathrm{d}p_a^*(y) \leq \liminf_{t\to\infty}\int_Y -\log\left(\frac{\tilde{r}_a(h_t)(y)}{\tilde{p}_a^*(y)}\right)\mathrm{d}p_a^{h_t}(y)$$

where the first inequality follows by Gibbs inequality and the second since by Assumption 2 (ii), $-\log\left(\frac{\tilde{r}_a(h_t)(y)}{\tilde{p}_a^*(y)}\right) \geq -K'$, p_a^*-almost surely and by Lemma 3.2 in Serfozo (1982). Therefore, we have

$$\lim_{t\to\infty}\sum_{a\in A}\sum_{\tau=1}^{t}\mathbb{I}_a(a_\tau)\int_Y -\log\left(\frac{\tilde{r}_a(h_t)(y)}{\tilde{p}_a^*(y)}\right)\mathrm{d}p_a^{h_t}(y) = 0.$$

So

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

Therefore the first part of the result follows by the continuity of the function defined in (10) and Theorem 17.31 in Aliprantis and Border (2013).

For the last part of the claim observe that

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

where the first equality follows from Lemma 2 and the third equality follows by the first part of the statement. ■

Define the set $\Theta^\varepsilon(a)$ as all parameters whose induced marginal distribution with respect to action a is at most ε away from a minimizer,

$$\Theta^\varepsilon(a) = \{\theta \in \Theta\colon \text{ there exists } \theta' \in \Theta(a) \text{ with } ||\theta - \theta'||_\infty \leq \varepsilon\}. \tag{11}$$

Lemma 6. *For every $a \in A$ and $\varepsilon \in (0,1)$, there exists $\epsilon \in (0,1)$ such that we have*

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

Proof. By Assumption 1 (ii) for every $y \in Y$, $\tilde{q}_a^{\hat{\theta}}(y)$ is a continuous function of $\hat{\theta}$. Therefore, since by Assumption 1 (iii) $\log\left(\frac{\tilde{q}_a^{\hat{\theta}}(y)}{\tilde{p}_a^*(y)}\right)$ is dominated by the function constant at K, by the Dominated Convergence Theorem

$$\hat{\theta} \mapsto \int_Y \log\left(\frac{\tilde{q}_a^{\hat{\theta}}(y)}{\tilde{p}_a^*(y)}\right) \mathrm{d}p_a^*(y) \tag{12}$$

is a continuous function. We now show that there exists $\epsilon \in (0,1)$ such that for all $\theta' \in \Theta \setminus \Theta^\varepsilon(a)$, for all $\theta \in \Theta^\epsilon(a)$

$$\int_Y \log\left(\frac{\tilde{q}_a^\theta(y)}{\tilde{q}_a^{\theta'}(y)}\right) \mathrm{d}p_a^*(y) = \int_Y \log\left(\frac{\tilde{q}_a^\theta(y)}{\tilde{p}_a^*(y)}\right) - \log\left(\frac{\tilde{q}_a^{\theta'}(y)}{\tilde{p}_a^*(y)}\right) \mathrm{d}p_a^*(y) > \epsilon.$$

Assume by contradiction that this does not hold true. Since Θ is compact, we would have 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

$$\int_Y \log\left(\frac{\tilde{q}_a^{\theta_n}(y)}{\tilde{p}_a^*(y)}\right) - \log\left(\frac{\tilde{q}_a^{\theta'_n}(y)}{\tilde{p}_a^*(y)}\right) \mathrm{d}p_a^*(y) \leq \frac{1}{n}.$$

But since the function of equation (12) is continuous, we would have

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

and

$$\int_Y \log\left(\frac{\tilde{q}_a^{\theta}(y)}{\tilde{p}_a^*(y)}\right) - \log\left(\frac{\tilde{q}_a^{\theta'}(y)}{\tilde{p}_a^*(y)}\right) \mathrm{d}p_a^*(y) = 0$$

a contradiction.

Then the result follows by the compactness of Θ and Assumption 1 (iii). ■

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

We first establish that if

$$\max_{\bar{a} \in A} \min_{y \in Y} u(\bar{a}, y) \leq \mathbb{E}_{p^*_{a'}}[u(a', y)] \qquad \forall a' \in A$$

then for every $c \in \mathbb{R}_{++}$ we have that $\beta_t = ct$ and is ε-safe. Let Π be a $\boldsymbol{\beta}$-optimal policy. Consider the stochastic process defined by

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

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 sequence (see Theorem 2.7 in Hall and Heyde, 2014 for the version that applies here), we have

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

so that

$$\liminf_{t\to\infty} \frac{\sum_{i=1}^{t} u(a_i, y_i)}{t} \geq \min_{a'\in A} \mathbb{E}_{p^*_{a'}}\left[u(a', y)\right] \geq \max_{\bar{a}\in A} \min_{y\in Y} u(\bar{a}, y) \qquad \mathbb{P}_\Pi\text{-a.s.}$$

Therefore, it is enough to check that the time scaling $\beta_t = ct$ that we propose below is such that every $\boldsymbol{\beta}$-optimal policy satisfies equation (3) for the $p^* \in \hat{P}$ with

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

For every $a \in \cup_{p^*\in\hat{P}} \underline{A}(p^*)$ let $\hat{p}^a \in \Delta(Y)$ be such that

$$\max_{\bar{a}\in A} \min_{y\in Y} u(\bar{a}, y) > \mathbb{E}_{\hat{p}^a_a}\left[u(a, y)\right]$$

and let $H := \max_{a\in\cup_{p^*\in\hat{P}}\underline{A}(p^*)} \min_{\theta\in\Theta} R\left(\hat{p}^a_a || q^\theta_a\right) < \infty$.

By the finiteness of $\underline{A}(p^*)$ there exists $\varphi^* > 0$ such that for every policy Π,

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

Observe that by definition of $\underline{A}(\cdot)$, there exists $J \in (0,1)$ such that for every exogenous $p^* \in \Delta(Y)^A$, $a' \in \underline{A}(p^*)$, and $\lambda \in \mathbb{R}_+$,

$$\mu\left(\left\{\theta : R\left(p^*_{a'} || q^\theta_{a'}\right) > J\right\}\right) \leq J \implies a' \notin BR^\lambda(\mu). \tag{13}$$

Let

$$c(p^*) := \frac{\varphi^*\left(\varepsilon + \max_{\bar{a}\in A}\min_{y\in Y} u(\bar{a}, y) - \max_{a'\in\underline{A}(p^*)} \mathbb{E}_{\hat{p}^{a'}_{a'}}\left[u(a', y)\right]\right)}{2HJ^2} \qquad \forall p^* \in \hat{P},$$

$c := \inf\left\{c(p^*) : p^* \in \hat{P}\right\} > 0$, and $\beta_t := ct$. Let Π be a $\boldsymbol{\beta}$-optimal policy.

Suppose by contradiction that there is a $\varphi \in (0,1)$ such that

$$\limsup_{t\to\infty} \alpha_t(a') = \varphi \geq \varphi^*$$

for some $\hat{a} \in \underline{A}(p^*)$ with positive $\mathbb{P}_\Pi$ probability. By Lemma 5 we have that

$$\begin{aligned}\liminf_{t\to\infty} \lambda^{\boldsymbol{\beta}}(h_t) &\geq \varphi \frac{\min_{\theta\in\Theta} R\left(p^*_{\hat{a}} || q^\theta_{\hat{a}}\right)}{2c} \geq \varphi \frac{J^2}{c} \\ &\geq \frac{2H}{\varepsilon + \max_{\bar{a}\in A} \min_{y\in Y} u(\bar{a}, y) - \max_{a'\in \underline{A}(p^*)} \mathbb{E}_{\hat{p}^{a'}_{a'}}[u(a', y)]} \qquad \mathbb{P}_\Pi\text{-a.s.}\end{aligned}$$

where the second inequality follows by equation (13). We show that if

$$\lambda \geq \frac{2H}{\varepsilon + \max_{\bar{a}\in A} \min_{y\in Y} u(\bar{a}, y) - \max_{a'\in \underline{A}(p^*)} \mathbb{E}_{\hat{p}^{a'}_{a'}}[u(a', y)]} \tag{14}$$

and

$$a \in BR^\lambda(\mu) \cap \left\{ a' \in A : \max_{\bar{a}\in A} \min_{y\in Y} u(\bar{a}, y) > \mathbb{E}_{p^*_{a'}}[u(a', y)] + \varepsilon \right\}$$

we have

$$\int_\Theta \min_{p_a \in \Delta(Y)} \mathbb{E}_{p_a}[u(a, y)] + \frac{1}{\lambda} R\left(p_a || q^\theta_a\right) d\mu(\theta) < \max_{a'\in A} \min_{p'\in\Delta(Y)} \mathbb{E}_{p'}[u(a', y)].$$

Indeed, observe that

$$\begin{aligned}&\int_\Theta \min_{p_a \in \Delta(Y)} \mathbb{E}_{p_a}[u(a, y)] + \frac{1}{\lambda} R\left(p_a || q^\theta_a\right) d\mu(\theta) \\ \leq\ & \int_\Theta \min_{p_a \in \Delta(Y)} \mathbb{E}_{p_a}[u(a, y)] \\ &+ \frac{R\left(p_a || q^\theta_a\right)\left(\varepsilon + \max_{\bar{a}\in A} \min_{y\in Y} u(\bar{a}, y) - \max_{a'\in \underline{A}(p^*)} \mathbb{E}_{p^*_{a'}}[u(a', y)]\right)}{2H} d\mu(\theta) \\ \leq\ & \int_\Theta \mathbb{E}_{\hat{p}^a_a}[u(a, y)] \\ &+ \frac{R\left(\hat{p}^a_a || q^\theta_a\right)\left(\varepsilon + \max_{\bar{a}\in A} \min_{y\in Y} u(\bar{a}, y) - \max_{a'\in \underline{A}(p^*)} \mathbb{E}_{\hat{p}^{a'}_{a'}}[u(a', y)]\right)}{2H} d\mu(\theta) \\ \leq\ & \mathbb{E}_{\hat{p}^a_a}[u(a, y)] + H \frac{\left(\varepsilon + \max_{\bar{a}\in A} \min_{y\in Y} u(\bar{a}, y) - \max_{a'\in \underline{A}(p^*)} \mathbb{E}_{\hat{p}^{a'}_{a'}}[u(a', y)]\right)}{2H} \\ \leq\ & \varepsilon + \max_{\bar{a}\in A} \min_{y\in Y} u(\bar{a}, y) + \max_{a'\in \underline{A}(p^*)} \mathbb{E}_{\hat{p}^{a'}_{a'}}[u(a', y)] < \max_{a'\in A} \min_{p'\in\Delta(Y)} \mathbb{E}_{p'}[u(a', y)]\end{aligned}$$

where the first inequality follows by equation (14) and the third inequality by the definition of H. But this is a contradiction with $a \in BR^\lambda(\mu)$, and so it cannot be played a positive fraction of periods. This in turn proves that $\boldsymbol{\beta}$ is ε-safe.

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

$$\min_{\theta \in \Theta} d\left(q_a^\theta, p^*\right) < \delta$$

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

$$\mathbb{E}_{p^*}\left[u\left(a', y\right)\right] - \mathbb{E}_{q_a^\theta}\left[u\left(a', y\right)\right] \leq \frac{\varepsilon}{4} \qquad \forall a' \in A.$$

But then $\beta = ct$ is ε-consistent with this δ by Lemma 6.

We next show that there is an exogenous decision problem (u, a, Y) such that no decision rule with $o\left(\beta_t\right) = t$ is $\frac{1}{2}$-safe. Suppose that

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

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

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

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

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

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

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

Therefore, by Lemma 2 we have that

$$\lim_{t \to \infty} \lambda^{\boldsymbol{\beta}}\left(h_t\right) = \lim_{t \to \infty} \frac{tR\left(0.6||0.4\right)}{\beta_t} = 0 \qquad \mathbb{P}_\Pi\text{-a.s.}$$

Moreover, for the constant function $\phi\left(\varepsilon\right) = \frac{1}{2}$ 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

$$\mu\left(0.4|h_t\right) \to 1 \qquad \mathbb{P}_{\Pi}\text{-a.s.}$$

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

$$\liminf_{t\to\infty} \frac{\sum_{i=1}^{t} u\left(a_i, y_i\right)}{t} = -0.2 < 0 = \max_{\bar{a}\in A} \min_{\bar{p}\in\Delta(Y)} \mathbb{E}_{\bar{p}}\left[u\left(\bar{a}, y\right)\right] \qquad \mathbb{P}_{\Pi}\text{-a.s.}$$

proving the desired result.

Finally, we show that there is an exogenous decision problem (u, a, Y) such that no decision rule with $\beta_t = o(t)$ is $\frac{1}{2}$-consistent. Let $\delta \in (0, 0.4)$ and suppose $A = \{1, -1, 0\}$, $Y = \{-1, 1\}$. 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 $\theta \in \Theta$

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

Let

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

For every $h_t \in C_t$, by Lemmas 2 and 5

$$\lambda^{\boldsymbol{\beta}}\left(h_t\right) \geq \frac{t^{\frac{R(0.6+\delta||0.6)}{2}}}{\beta_t}$$

that by definition is diverging to $+\infty$ and so 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. ■

The next lemma simplifies the task of proving that an action b cannot be a $\boldsymbol{\beta}$-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 strategy must be preformed with respect to all priors

supported over Θ and "time-shifts" of the actual concern for misspecification.

Lemma 7. *Let $b \in A$. If $\mathbb{P}_{\Pi^b}[b \in BR^{\frac{k+LRT(h_\tau,\Theta)}{2\beta_{\tau+n}}}(\nu(\cdot|h_\tau))$ for all $\tau \geq 0] = 0$ for every prior belief ν with support Θ, every $k \in \mathbb{R}_+$ and $n \in \mathbb{N}_0$, then b is not a $\boldsymbol{\beta}$-limit action.*

Proof. Suppose by way of contradiction that there are a $\boldsymbol{\beta}$-optimal policy $\tilde{\Pi}$ and a history (a^t, y^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

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

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

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

However, the latter equals zero by the assumption of the lemma, which establishes that b cannot be a $\boldsymbol{\beta}$-limit action. ■

Proof of Proposition 2. Suppose that a^* is not a self-confirming equilibrium. By Lemma 7 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 strong law of large numbers (see, e.g., Theorem 8.3.5 in Dudley, 2018),

$$\lim_{t\to\infty} p^{h_t}_{a^*} = p^*_{a^*} \qquad \mathbb{P}_{\Pi_{a^*}}\text{-a.s.}$$

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

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

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

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

With this, as by Lemma 6 for every $\varepsilon \in \mathbb{R}_{++}$ we have

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

the desired conclusion follows by Lemma 4. ■

Lemma 8. *For every $a \in A$ and $\theta \in \Theta$, there exists a $\mathcal{B}(Y)$ measurable Radon-Nykodim derivative between q_a^θ and ξ_a.*

Proof. Since q_a^θ is a probability measure, it has bounded variation and therefore the lemma follows directly by Theorem 13.18 in Aliprantis and Border (2013). ■

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

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

Proof. Let $\hat{y} \in \text{argmin}_{y\in\cup_{\theta\in\Theta}\text{supp}q_a^\theta} u(a,y)$. 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

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

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

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

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

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

where the equality follows by 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 by Lemma 6 in Dunford and Schwartz (1988). ■

Proof of Theorem 1. 1) Suppose that a^* is not a Berk-Nash equilibrium. By Lemma 7 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 $Y \subseteq \mathbb{R}^k$), $p_{a^*}^{h_t} \to p_{a^*}^*$, $\mathbb{P}_{\Pi_{a^*}}$-a.s. Moreover, $\tilde{p}_{a_t}^*(y_t) > 0$ for all $t \in \mathbb{N}$, $\mathbb{P}_{\Pi_{a^*}}$-almost surely.

Then, by Lemma 5, since $t = o(\beta_t)$, we have

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

By Assumption 1 and Lemma 8, the assumptions of Berk (1966), page 54, are satisfied, and we have that for every $\varepsilon \in \mathbb{R}_{++}$,

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

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 $\nu^* \in \Delta\left(\Theta(a^*)\right)$. With this, the result follows by 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 7 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 $Y \subseteq \mathbb{R}^k$),

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

Moreover, $\tilde{p}_{a_t}^*(y_t) > 0$ for all $t \in \mathbb{N}$, $\mathbb{P}_{\Pi_{a^*}}$-almost surely.

Then, by Lemma 5, since $\beta_t = o(t)$ and $\min_{\theta\in\Theta} R\left(p_a^* || q_{a^*}^\theta\right) > 0$, we have

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

Moreover, we have

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

where the inequality follows from the first part of Lemma 5. 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

$$\lim_{t\to\infty} \inf_{\theta\in\Theta} \min_{p_a \in \Delta(Y)} \int_Y u(a,y)\, \mathrm{d}p_a + \frac{R\left(p_a || q_a^\theta\right)}{\frac{k+LRT(h_t,\Theta)}{2\beta_{n+t}}} = \min_{y \in \cup_{\theta\in\Theta} \mathrm{supp} q_a^\theta} u(a,y).$$

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

$$\lim_{t\to\infty} \mathbb{E}_{\mu(\cdot|h_t)} \left[\min_{p_a \in \Delta(Y)} \int_Y u(a,y)\, \mathrm{d}p_a + \frac{R\left(p_a || q_a^\theta\right)}{\lambda^{\boldsymbol{\beta}}(h_t)} \right] = \min_{y \in \cup_{\theta\in\Theta} \mathrm{supp} q_a^\theta} u(a,y) \quad \mathbb{P}_{\Pi_{a^*}}\text{-a.s.}$$

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

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

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

we have $a^* \notin BR^{R\left(p_a^* || q_a^{\hat{\theta}}\right)/c}(\nu)$. By Lemma 7 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

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

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 $Y \subseteq \mathbb{R}^k$),

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

Moreover, $\tilde{p}_{a_t}^*(y_t) > 0$ for all $t \in \mathbb{N}$, $\mathbb{P}_{\Pi_{a^*}}$-almost surely.

By Lemma 5

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

By Assumption 1 and Lemma 8, the assumptions of Berk (1966), page 54, are satisfied, and we have that for every $\varepsilon \in \mathbb{R}_{++}$, $\nu\left(\Theta_\varepsilon(a^*)|h_t\right) \to 1$, $\mathbb{P}_{\Pi_{a^*}}$-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 $\left(\nu\left(\cdot|h_t\right)\right)_{t\in\mathbb{N}}$ is converging to a $\bar{\nu} \in \Delta(\Theta)$ with

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

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

Lemma 10. *For every $c \in \mathbb{R}_{++}$ the function $\alpha \mapsto \min_{\theta\in\Theta} \sum_{a\in A} \alpha(a) R\left(p_a^* || q_a^\theta\right)/c$ is*

continuous and the correspondence $\Theta(\cdot) : \Delta(A) \to 2^{\Theta}$ *is upper hemicontinuous.*

Proof. We first show that the function

$$(\alpha, \theta) \mapsto \sum_{a \in A} \alpha(a) R\left(p_a^* || q_a^{\theta}\right) / c \tag{15}$$

is continuous. Fix an $a \in A$ and let $(\theta_n)_{n \in \mathbb{N}}$ be a sequence that converges to θ. By Theorem 1 in Sweeting (1986) and Assumption 1 (ii), $\tilde{q}_a^{\theta_n}$ is uniformly converging to $\tilde{q}_a^{\theta}$. Then

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

and observe that the integrand on the right hand side is dominated by $2K$ by Assumption 1 (iii). Therefore, by the dominated convergence theorem $|R\left(p_a^* || q_a^{\theta_n}\right) - R\left(p_a^* || q_a^{\theta}\right)|$ converges to 0. Since A is finite and the function in equation (15) is linear in α, we have obtained the desired continuity. With this, the statement follows by 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

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

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 $\left[0, \bar{\lambda}\right]$ where

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

where the inequality holds by Assumption 1 (iii). 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 their first and third argument and U_3 is continuous in its first and secund argument by Lemma 10. 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

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

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

and

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

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

Proof of Theorem 2. Observe that $(\alpha_t)_{t\in\mathbb{N}}$ satisfies the following differential inclusion: for all $a \in A$, $t \in \mathbb{N}$, $h_t \in H_t$, and $h_{t+1} \in H_{t+1}$ such that $h_{t+1} \succ h_t$

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

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 $w : \mathbb{R}_+ \to \Delta(A)$

$$w(\tau_t + l) = \alpha_t + l\frac{\alpha_{t+1} - \alpha_t}{\tau_{t+1} - \tau_t}, \quad l \in \left[0, \frac{1}{t+1}\right]. \tag{16}$$

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

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

We will use the theory of stochastic approximation for differential inclusions (Benaim, Hofbauer, and Sorin, 2005) to show that (16) can be approximated by a solution to

$$\dot{\boldsymbol{\alpha}}_t \in \chi_{\boldsymbol{\alpha}_t} - \boldsymbol{\alpha}_t. \tag{17}$$

A solution over $[0,T]$, $T \in \mathbb{R}_{++}$, to the differential inclusion (17) with initial point $\hat{\alpha} \in \Delta(A)$ is a mapping $l : [0,T] \to \Delta(A)$ that is absolutely continuous over compact intervals such that $l(0) = \hat{\alpha}$ and (17) is satisfied for almost every t. Let $L_{\hat{\alpha}}^T$ be the set of the solutions to (17) over $[0,T]$, $T \in \mathbb{R}_{++}$, with initial conditions $\hat{\alpha} \in \Delta(A)$. Since by Lemmas 10 and 4 $\alpha \mapsto \min_{\theta\in\Theta}\sum_{a\in A}\alpha\left(a\right)R\left(p_a^*||q_a^\theta\right)/c$ is continuous and $\Theta\left(\cdot\right)$, $BR^{(\cdot)}\left(\cdot\right)$ are upper hemicontinuous,

$$\alpha \mapsto \Delta\left(BR^{\min_{\theta\in\Theta}\sum_{a\in A}\alpha(a)R\left(p_a^*||q_a^\theta\right)/c}\left(\Delta\left(\Theta\left(\alpha\right)\right)\right)\right) \tag{18}$$

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

$$\left(\alpha_n, \alpha_n'\right) \in \left(\Delta\left(A\right) \times \Delta\left(A\right)\right)^{\mathbb{N}}$$

be such that

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

and $\lim_{n\to\infty}\left(\alpha_n, \alpha_n'\right) = \left(\alpha, \alpha'\right)$. By finiteness of the actions, we can take $\left(\alpha_n, \alpha_n'\right)_{n\in\mathbb{N}}$ 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 a $\left(\nu_n^{\bar{a}}\right)_{n\in\mathbb{N}}$ such $\nu_n^{\bar{a}} \in \Delta\left(\Theta\left(\alpha_n'\right)\right)$ and

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

Since Θ is compact so is $\Delta\left(\Theta\right)$ 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 $\nu^{\bar{a}} \in \Delta\left(\Theta\right)$.

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

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

showing that (α, α') belongs to the graph of correspondence (18). Therefore a solution to (17) exists by Theorem 2.1.4 in Aubin and Cellina (2012), i.e., $L_{\hat{\alpha}}^T$ is nonempty for every $T \in \mathbb{R}_{++}$ and $\hat{\alpha} \in \Delta(A)$.

We next establish that the continuous-time interpolation of $(\alpha_t(h_t))_{t\in\mathbb{N}}$ defined in (16) can in the long run be approximated arbitrarily well by a solution to (17). Observe that w is Lipschitz continuous of order 1 as for all $t \in \mathbb{N}$,

$$\frac{||\alpha_{t+1} - \alpha_t||_\infty}{\tau_{t+1} - \tau_t} \leq \frac{1/(t+1)}{1/(t+1)} = 1.$$

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 $\varepsilon \in \mathbb{R}_+$

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

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

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

by

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

Observe that $F(0, 0, \alpha) = \chi_\alpha - \alpha$. Moreover, we now show that F has a closed graph, and 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

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

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 of $\left(\nu_n^{\hat{a}}, \lambda_n^{\hat{a}}\right)_{n\in\mathbb{N}}$ such that

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

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

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

is continuous by Lemma 10,

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

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

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

showing that $(\iota, \varepsilon, \varepsilon', \alpha)$ belongs to the graph of the correspondence (18).

With this, $F(0, 0, \cdot) + (\cdot) : \Delta(A) \to \Delta(A)$ satisfies Hypothesis 1.1 in Benaim, Hofbauer, and Sorin (2005), as it is clearly compact- and convex-valued. Moreover, since $\tilde{p}_{a_t}^*(y_t) > 0$ for all $t \in \mathbb{N}$, $\mathbb{P}_\Pi$-almost surely, by Theorem 1 in Esponda, Pouzo, and Yamamoto (2021) and Lemma 5 we have that $\mathbb{P}_\Pi$-almost surely, if $\lim_{t\to\infty}\alpha_t = \alpha^*$, we eventually have $\mu(\cdot|h_t) \in M_\varepsilon(\alpha^*)$ and

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

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 of $\chi_{\alpha_t} - \alpha_t \in F(\hat{\varepsilon}_t, \hat{\varepsilon}_t, \alpha_t)$. By Proposition 1.3 in Benaim, Hofbauer, and Sorin (2005) w

is what they call a perturbed solution of (17) and by their Theorem 4.2

$$\lim_{t\to\infty} \inf_{\tilde{\boldsymbol{\alpha}}\in L^T_{w(t)}} \sup_{0\le s\le T} ||w(t+s) - \tilde{\boldsymbol{\alpha}}_s|| = 0 \quad \mathbb{P}_\Pi\text{-a.s. for all } T \in \mathbb{N}. \tag{19}$$

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 (17), $\boldsymbol{\alpha}(a)$ 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 (19) 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 $t \ge \hat{T}$ and

$$\inf_{\tilde{\boldsymbol{\alpha}}\in L^{T^*}_{w(\hat{T})}} \sup_{0\le s\le T^*} ||w(\hat{T}+s) - \tilde{\boldsymbol{\alpha}}_s|| \le D/2. \tag{20}$$

But then by definition of T^* the differential inclusion leaves $B_D(\alpha^*)$ by time $T^* + \hat{T}$, and by (20), α_t does not stay in $B_{D/2}(\alpha^*)$, a contradiction. ■

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

$$\Theta(\alpha) = (\theta_0^*, \theta_{1\pi}^*, \theta_{1a}^*, \theta_2^*, \theta_3^*). \tag{21}$$

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

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

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 $\min_{\theta\in\Theta} R\left(p_1^*||q_1^\theta\right) > 0$. By Assumption 3 (ii) and Lemma 9 there exists a sufficiently small $\bar{c}$ such that for all $c \leq \bar{c}$,

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

proving that there is no c-robust equilibrium if $c \leq \bar{c}$. That a mixed c-robust equilibrium exists follows by Proposition 3.[33]

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

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 K is denoted as $B_0(\Sigma, K)$ (resp. $B(\Sigma, K)$). A functional $I : \Phi \to \mathbb{R}$ defined on a nonempty subset Φ of $B(\Sigma)$ is a *niveloid* if it is monotone and 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 $k \in K$. 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 $c : \Delta(S) \to \mathbb{R}_+$ 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 $x \succ x'Ex$. If $f \in \mathcal{F}$, an element $x_f \in X$ is a certain equivalent of f if $f \sim x_f$. An act $f \in \mathcal{F}$ is *structured* if it is measurable with respect to Σ_s.

Definition 13. A tuple (u, Q, μ, ϕ) is an *average second order utility representation*

[33] To 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 (21) every c-robust ergodic equilibrium with the reduced parameter space remains a c-robust ergodic equilibrium with the original Θ.

of the preference relation $\succsim$ if $u : X \to \mathbb{R}$ is a nonconstant affine function, $Q \subseteq \Delta(S)$ is a finite nonempty set, $\mu \in \Delta(Q)$, $\phi : u(X) \to \mathbb{R}$ is a strictly increasing continuous function, and

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

A.2.2 Results

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

Lemma 11. *Suppose that there exist a nonconstant affine function $u : X \to \mathbb{R}$, 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 $f, g \in \mathcal{F}$*

$$\begin{array}{ccc} f & \succsim & g \\ & \Longleftrightarrow & \\ \mathbb{E}_\mu\left[\min\limits_{p\in\Delta(S)} \int_S u(f)\,dp + \frac{1}{\lambda_q} R(p||q)\right] & \geq & \mathbb{E}_\mu\left[\min\limits_{p\in\Delta(S)} \int_S u(g)\,dp + \frac{1}{\lambda_q} R(p||q)\right]. \end{array}$$

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

$$f \succsim g \Longleftrightarrow \min_{p\in\Delta(S)} \int_S \hat{u}(f)\,dp + \hat{c}(p) \geq \min_{p\in\Delta(S)} \int_S \hat{u}(g)\,dp + \hat{c}(p) \tag{22}$$

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

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

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

is easily seen to be monotone, translation invariant, and concave. Since Q is finite,

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

is well-defined. By Theorem 11.5 in Aliprantis and Border (2013), $\hat{I}$ is monotone and concave. Let $\varphi \in B_0(\Sigma, u(X))$, $k \in u(X)$, and $\gamma \in (0,1)$. We have

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

where the second equality follows from the translation invariance of each I_q, and the third equality follows by Theorem 11.5 in Aliprantis and Border (2013). But then, notice that

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

where $\hat{I}$ is monotone and translation invariant. Therefore, by Lemma 25 in Maccheroni, Marinacci, and Rustichini (2006a), $\hat{I} : B_0(\Sigma, u(X)) \to \mathbb{R}$ is a concave niveloid. 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 1 and (iii) $\Longrightarrow$ (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 $E_1 \supseteq E_2 \supseteq \ldots$, $\cap_{n\geq 1} E_n = \emptyset$, 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) $\Longrightarrow$ (ii) of Theorem 50 in Cerreia-Vioglio, Maccheroni, Marinacci, and Montrucchio (2011) there exists $i \in \mathbb{N}$ such that $\hat{I}\left(k\mathbb{I}_{E_i} + \varphi\mathbb{I}_{S\setminus E_i}\right) > \hat{I}(\varphi')$. But observe that

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

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 (22).

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 Theorem 3 in Maccheroni, Marinacci, and Rustichini (2006a) and the fact that each two affine functions that represent $\succsim$ on X are positive affine transformations of each other (Herstein and Milnor, 1953), we can choose $u = \hat{u}$. Finally, by Lemma 32 in Maccheroni, Marinacci, and Rustichini (2006a) for every $q \in Q$, and $k \in u(X)$, $\partial I_q(k) = \{q\}$. Let $\bar{k} \in \operatorname{int} u(X) \neq \emptyset$ and observe that since Q is finite,

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

Applying Lemma 32 in Maccheroni, Marinacci, and Rustichini (2006a) again we obtain that the unique $\hat{c}$ identified by the choice of $\hat{u}$ has

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

■

Lemma 12. *Let $E \in \Sigma_{st}$ be nonnull, and let $\succsim$ satisfy 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 $f, g \in \mathcal{F}$, $x \in X$, $(E_i)_{i\in\mathbb{N}} \in \Sigma^{\mathbb{N}}$ with $E_1 \supseteq E_2 \supseteq \ldots$ and $\cap_{n\geq 1} E_n = \emptyset$, and $f \succ_E g$. Therefore, $(E'_i)_{i\in\mathbb{N}}$ where $E'_i = E_i \cap E$ is such that $E'_1 \supseteq E'_2 \supseteq \ldots$ 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 $\left(xE'_{n_0} f\right) Eh \succ gEh$. But notice that

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

and therefore $xE_{n_0} f \succ_E g$. ■

Lemma 13. *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, Weak Monotone Continuity, and the Intramodel Sure-Thing Principle. For every $f, g \in \mathcal{F}$, we have*

$$\begin{array}{ccc} f & \succsim_\rho & g \\ & \Longleftrightarrow & \\ \min\limits_{q\in\Delta(S)} \displaystyle\int u_\rho\left(f\left(s\right)\right) dq\left(s\right) + \frac{1}{\lambda_\rho} R\left(q||p_\rho\right) & \geq & \min\limits_{q\in\Delta(S)} \displaystyle\int u_\rho\left(f\right) dq + \frac{1}{\lambda_\rho} R\left(q||p_\rho\right) \end{array} \tag{23}$$

where u_ρ is a nonconstant affine function, $\lambda_\rho \in [0, \infty)$, and $p_\rho \in \Delta\left(S\right)$. 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 $\operatorname{supp} p_\rho \subseteq \Omega \times \{\rho\}$.

Proof. By Lemma 12 $\succsim_\rho$ satisfies Weak Order, Weak Certainty Independence, Continuity, Monotonicity, Uncertainty Aversion, and Weak Monotone Continuity. 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 $\succsim_\rho$, $fEh \succsim_\rho gEh$ implies that there exists $\hat{h} \in \mathcal{F}$ such that

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

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

$$\begin{aligned}
& (fEh)\,\rho\hat{h} \succsim (gEh)\,\rho\hat{h} \\
\Longrightarrow\quad & (f\{(\omega,\rho):(\omega,\rho)\in E\}\,h)\,\rho\hat{h} \succsim (f\{(\omega,\rho):(\omega,\rho)\in E\}\,h)\,\rho\hat{h} \\
\Longrightarrow\quad & (f\{(\omega,\rho'):\rho'\in\Delta(\Omega),(\omega,\rho)\in E\}\,h)\,\rho\hat{h} \succsim (f\{(\omega,\rho'):\rho'\in\Delta(\Omega),(\omega,\rho)\in E\}\,h)\,\rho\hat{h} \\
\Longrightarrow\quad & (f\{(\omega,\rho'):\rho'\in\Delta(\Omega),(\omega,\rho)\in E\}\,h) \succsim_\rho (f\{(\omega,\rho'):\rho'\in\Delta(\Omega),(\omega,\rho)\in E\}\,h) \\
\Longrightarrow\quad & \left(f\{(\omega,\rho'):\rho'\in\Delta(\Omega),(\omega,\rho)\in E\}\,\bar{h}\right) \succsim_\rho \left(f\{(\omega,\rho'):\rho'\in\Delta(\Omega),(\omega,\rho)\in E\}\,\bar{h}\right) \\
\Longrightarrow\quad & \left(f\{(\omega,\rho'):\rho'\in\Delta(\Omega),(\omega,\rho)\in E\}\,\bar{h}\right)\rho h' \succsim \left(f\{(\omega,\rho'):\rho'\in\Delta(\Omega),(\omega,\rho)\in E\}\,\bar{h}\right)\rho\hat{h} \\
\Longrightarrow\quad & \left(f\{(\omega,\rho):(\omega,\rho)\in E\}\,\bar{h}\right)\rho h' \succsim \left(f\{(\omega,\rho):(\omega,\rho)\in E\}\,\bar{h}\right)\rho\hat{h} \\
\Longrightarrow\quad & \left(fE\bar{h}\right)\rho h' \succsim \left(gE\bar{h}\right)\rho h' \implies fE\bar{h} \succsim_\rho gE\bar{h}
\end{aligned}$$

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

With this, the first part follows by 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:

$$f \succsim g \Longleftrightarrow \min_{p\in\Delta(S)}\left(\int u(f)\,dp + c(p)\right) \geq \min_{p\in\Delta(S)}\left(\int u(g)\,dp + c(p)\right) \tag{24}$$

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

Next, assume $\Omega \times \{\rho\}$ is strongly nonnull. Notice that $\succsim$ and $\succsim_\rho$ 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 (24) $u(x) \geq u(x')$. Since c is grounded, there exists $q^* \in \Delta(S)$ with $c(q^*) = 0$. But then

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

that is, $x\,(\Omega\times\{\rho\})\,x \succsim x'\,(\Omega\times\{\rho\})\,x$, and $x \succsim_\rho x'$. Therefore, by the mixture space theorem (Herstein and Milnor, 1953), u and u_ρ are positive affine transformations of each other, proving that u_ρ can be chosen to be the same for all such ρ. 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,

$$\begin{aligned} u(x) &> u(y)\, p_\rho(E) + u(x)\, p_\rho(E) \\ &\geq \min_{q \in \Delta(S)} \int u(y)\, q(E) + u(x)\, q(E)\, dq + \frac{1}{\lambda_\rho} R(q||p_\rho) \end{aligned}$$

where the second inequality follows by equation (23), and so $x \succ_\rho yEx$. But since $x = x(\Omega \times \{\rho\})\, x$, $yEx = x(\Omega \times \{\rho\})(yEx)$ and $\Omega \times \{\rho\} \in \Sigma_{st}$ this would imply $x \succ_\rho x$, a contradiction by Lemma 12. ■

Lemma 14. *(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 $q \in \Delta(S)$

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

Therefore,

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

and the result follows by 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 $q \in \Delta(S)$

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

Therefore,

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

■

Lemma 15. *Suppose that the assumptions of Theorem 3 hold. Let* $\succsim$ *be such that*

$$\begin{array}{ccc} f & \succsim & g \\ & \Longleftrightarrow & \\ \mathbb{E}_\mu\left[\min_{p\in\Delta(S)}\int_S u(f)\,dp+\frac{1}{\lambda_q}R(p||q)\right] & \geq & \mathbb{E}_\mu\left[\min_{p\in\Delta(S)}\int_S u(g)\,dp+\frac{1}{\lambda_q}R(p||q)\right] \end{array}$$

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

$$\begin{array}{ll} Q & \to \Delta(\Omega) \\ q & \mapsto \rho_q \end{array}$$

with

$$q\left(\left\{\omega, \rho_q\right\}\right) = \rho_q(\omega) \qquad \forall q \in Q, \forall \omega \in \Omega,$$

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

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

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

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

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

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

2) It immediately follows from 1). ■

Lemma 16. *Suppose that the assumptions of Theorem 3 hold. Let* $\succsim$ *be such that*

$$\begin{aligned}
f & \succsim g \\
& \Longleftrightarrow \\
\mathbb{E}_\mu\left[\min_{p\in\Delta(S)} \int_S u\left(f\right) dp + \frac{1}{\lambda_q} R\left(p||q\right)\right] & \geq \mathbb{E}_\mu\left[\min_{p\in\Delta(S)} \int_S u\left(g\right) dp + \frac{1}{\lambda_q} R\left(p||q\right)\right]
\end{aligned}$$

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

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

with

$$q\left(\left\{\omega, \rho_q\right\}\right) = \rho_q\left(\omega\right) \qquad \forall q \in Q, \forall \omega \in \Omega$$

$\mu \in \Delta(Q)$, 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 of Lemma 16. (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,

$$\rho(\{\omega : f(\omega, \rho) = y\}) = \rho'(\{\omega : g(\omega, \rho') = y\}) \quad \forall y \in X, \tag{25}$$

and $f \succsim_{\Omega\times\{\rho\}} x$. Since $\Omega \times \{\rho\}$ and $\Omega \times \{\rho'\}$ are nonnull, by part 2 of Lemma 15 there exist $q, q' \in Q$ with $\mu(q) > 0$, $\mu(q') > 0$, $\rho_q = \rho$, and $\rho_{q'} = \rho'$. Let

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

and let $\xi \in \Delta(X)$ be the finite support probability measure such that

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

Then, equation (25) implies

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

Therefore, by Lemma 15 both $f \succsim_{\Omega\times\{\rho\}} x$ and $f \succsim_{\Omega\times\{\rho'\}} x$ mean that

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

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

$$\lambda_q > k > \lambda_{q'}. \tag{26}$$

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

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

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

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

for all $x \in X$. By the convexity of X and Lemma 15 there exists $\hat{x} \in X$ with $z \succ \hat{x} \succ y$ and

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

But by equation (26) and Lemma 15 we have

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

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 by Lemma 11.

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 15, 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

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

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

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

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

$$\begin{aligned} & \int_{W\times\Delta(\Omega)} \phi(u(f))\,\mathrm{d}\bar{q} + \int_{(\Omega\setminus W)\times\Delta(\Omega)} \phi(u(h))\,\mathrm{d}\bar{q} \\ \geq & \int_{W\times\Delta(\Omega)} \phi(u(g))\,\mathrm{d}\bar{q} + \int_{(\Omega\setminus W)\times\Delta(\Omega)} \phi(u(h))\,\mathrm{d}\bar{q}. \end{aligned} \tag{27}$$

or

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

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 16. 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 15. Let $\Omega \times B \in \Sigma_s$ and $f, g, h, h' \in \mathcal{F}$. If either $\Omega \times B$ or $\Omega \times (\Delta(\Omega) \setminus B)$ is null, we clearly have that $\Omega \times B \in \Sigma_{st}$. Suppose $\Omega \times B$ and $\Omega \times (\Delta(\Omega) \setminus B)$ are nonull, then

$$\begin{aligned} fEh & \succsim gEh \\ & \Longleftrightarrow \\ \int_{\{q:\rho_q\in B\}} \min_{p\in\Delta(S)} \int_S u(f)\,\mathrm{d}p + \frac{1}{\lambda_q} R(p||q)\,\mathrm{d}\mu(q) & \geq \int_{\{q:\rho_q\in B\}} \min_{p\in\Delta(S)} \int_S u(g)\,\mathrm{d}p + \frac{1}{\lambda_q} R(p||q)\,\mathrm{d}\mu(q) \\ & \Longleftrightarrow \\ fEh' & \succsim gEh' \end{aligned}$$

where the two equivalences follow by Lemma 15. Moreover

$$\begin{aligned} hEf & \succsim hEg \\ & \Longleftrightarrow \\ \int_{\{q:\rho_q\notin B\}} \min_{p\in\Delta(S)} \int_S u(f)\,\mathrm{d}p + \frac{1}{\lambda_q} R(p||q)\,\mathrm{d}\mu(q) & \geq \int_{\{q:\rho_q\notin B\}} \min_{p\in\Delta(S)} \int_S u(g)\,\mathrm{d}p + \frac{1}{\lambda_q} R(p||q)\,\mathrm{d}\mu(q) \\ & \Longleftrightarrow \\ h'Ef & \succsim h'Eg \end{aligned}$$

where again the two equivalences follow by Lemma 15. 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 15 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 $\bar{x}, \underline{x} \in X$, as well as $\bar{s}, \underline{s} \in E$ with $f(\bar{s}) = \bar{x}$ and $g(\underline{s}) = \underline{x}$ and

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

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, $\bar{x}\emptyset\bar{x} = \bar{x}E\bar{x} \succsim fE\bar{x} \succ gE\bar{x} \succsim \underline{x}E\bar{x}$. Therefore, by Structured Savage P4, $x = x\emptyset x \succ x'Ex$. Since $E \in \Sigma_s$, $x, x' \in X$ 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 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 13 we have

$$f \succsim_\rho g \iff \min_{p \in \Delta(S)} \int_S \hat{u}(f)\, \mathrm{d}p + \frac{1}{\lambda_\rho} R(p||q_\rho) \tag{28}$$

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

$$Q = \{q_\rho \in \Delta(S) : \rho \in B\}. \tag{29}$$

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

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

where λ_{ρ_q} is given by equation (28).

We now show that

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

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 1 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, and Theorem III.6.6 in Wakker (2013) for every $q \in Q$ there exists strictly increasing $V_q : \hat{u}(X) \to \mathbb{R}$ such that for all $\psi, \psi' \in \hat{u}(X)^Q$

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

By Structured Savage P2, $\left\{\Omega \times \rho_q\right\}_{q \in Q} \subseteq \Sigma_{st}$, and therefore we have $V_q(l) = al + b$ for some $a \in \mathbb{R}_{++}$ and $b \in \mathbb{R}_m$, so that we can rewrite the previous equivalence as

$$\psi \succsim \psi' \iff \sum_{q \in Q} V(\psi(q))\,\mu(q)$$

for some strictly increasing $V : \hat{u}(X) \to \mathbb{R}$ and $\mu \in \Delta(Q)$.

Moreover, by Uncertainty Neutrality over Models and since the state space is Q-adequate, it follows that V is midpoint affine. In addition, it is measurable since it is increasing. Hence, by Sierpinski's theorem (see, e.g., Theorem 72C in Roberts and Varberg (1973)) it is affine. Therefore, as the representation is unique up to a positive linear transformation (see, Observation III.6.6 Wakker (2013)), we can take V equal to the identity.

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

$$\begin{array}{ccc} f & \succsim & g \\ & \iff & \\ \mathbb{E}_{\hat{\mu}}\left[\min_{p \in \Delta(S)} \int_S u(f)\,\mathrm{d}p + \frac{1}{\lambda_q} R(p||q)\right] & \geq & \mathbb{E}_{\hat{\mu}}\left[\min_{p \in \Delta(S)} \int_S u(g)\,\mathrm{d}p + \frac{1}{\lambda_q} R(p||q)\right] \end{array}$$

for some some probability measure $\hat{\mu} \in \Delta(Q)$.

Moreover, by Lemma 16 Uniform Misspecification Concern implies that $\lambda = \lambda_E$ for all $E \in \Sigma_s$, proving the result. ■

Proof of Corollary 2. It immediately follows from Lemma 11 and Proposition 8 in Maccheroni, Marinacci, and Rustichini (2006a). ■

Proof of Proposition 4. That each $\succsim^h$ admits an average robust control representation $(u_h, Q_h, \mu_h, \lambda_h)$ follows by Theorem 3.[34] 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 by induction on t 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 $\succsim^{\emptyset}_{\omega_i}$ is nontrivial for all $i \in \{1, ..., t\}$. By definition we have $\mu_{h_t} = \mu$ for the empty history. Suppose that the statement holds for all the histories of length $t \in \mathbb{N}_0$, and let $h_{t+1} = (\omega_i)_{i=1}^{t+1} = (h_t, \omega_{t+1})$ be an history such that $\succsim^{\emptyset}_{\omega_i}$ is nontrivial for all $i \in \{1, ..., t+1\}$. Then $h_t = (\omega_i)_{i=1}^t$ is an history such that $\succsim^{\emptyset}_{\omega_i}$ is nontrivial for all $i \in \{1, ..., t\}$ and by the inductive hypothesis $\succsim^{h_t}$ admits an average robust control representation $(u, Q_{h_t}, \mu(\cdot|h_t), \lambda_{h_t})$. Let f and g be measurable with respect to Σ_s. Then

$$f \succsim^{h_t}_{\{\omega_{t+1}\} \times \Delta(\Omega)} g$$

means that

$$\begin{aligned} \int_{\Delta(\Delta(S))} \rho_q(\omega_{t+1}) u\left(f\left(\omega_{t+1}, \rho_q\right)\right) d\mu(q|h_t) &\geq \int_{\Delta(\Delta(S))} \rho_q(\omega_{t+1}) u\left(g\left(\omega_{t+1}, \rho_q\right)\right) d\mu(q|h_t) \\ \iff \\ \frac{\int_{\Delta(\Delta(S))} \rho_q(\omega_{t+1}) u\left(f\left(\omega_{t+1}, \rho_q\right)\right) d\mu(q|h_t)}{\int_{\Delta(\Delta(S))} \rho_q(\omega_{t+1}) d\mu(q|h_t)} &\geq \frac{\int_{\Delta(\Delta(S))} \rho_q(\omega_{t+1}) u\left(g\left(\omega_{t+1}, \rho_q\right)\right) d\mu(q|h_t)}{\int_{\Delta(\Delta(S))} \rho_q(\omega_{t+1}) d\mu(q|h_t)}. \end{aligned}$$

Since f and g are measurable with respect to Σ_s we can suppress their dependence on

[34] To avoid long technical detours and invoke Theorem 3 the statement maintains the assumption that the state space remains Q-adequate at every period. This is unnecessarily restrictive, as often some sequences of observations may induce updated beliefs that do not have ties. It is straightforward to adapt the arguments of the theorem under the weaker form of Uncertainty Neutrality over Models that for every conditional preference mixes with the updated likelihood ratio between model. This issue disappears at all when the agent has a nonatomic belief over models.

ω_{t+1} to obtain:

$$\int_{\Delta(\Delta(S))} u\left(f\left(\rho_q\right)\right) d\mu\left(q|\left(h_t,\omega_{t+1}\right)\right) \geq \int_{\Delta(\Delta(S))} u\left(g\left(\rho_q\right)\right) d\mu\left(q|\left(h_t,\omega_{t+1}\right)\right)$$

that is, $\succsim^{h_t}_{\{\omega_{t+1}\}\times\Delta(\Omega)}$ admits a SEU representation of the acts measurable with respect to Σ_s with Bernoulli utility u and probability measure $\mu\left(\cdot|\left(h_t,\omega_{t+1}\right)\right)$. Then, by the equality $\succsim^{h_t}_{\{\omega_{t+1}\}\times\Delta(\Omega)}=\succsim^{(h_t,\omega_{t+1})}$ imposed by Dynamic Consistency over Models, $\succsim^{(h_t,\omega_{t+1})}$ also admits an average robust control representation

$$\left(u, Q_{(h_t,\omega_{t+1})}, \mu\left(\cdot|\left(h_t,\omega_{t+1}\right)\right), \lambda_{(h_t,\omega_{t+1})}\right)$$

that is SEU for the acts measurable with respect to Σ_s with Bernoulli utility u and probability measure $\mu\left(\cdot|\left(h_t,\omega_{t+1}\right)\right)$. Finally that $Q_{h_t} = Q$ follows from the fact that since $\left\{\rho_q : q \in Q\right\} \subseteq \operatorname{int}\Delta\left(\Omega\right)$, we have $\operatorname{supp}\mu\left(\cdot|\left(h_t\right)\right) \subseteq \operatorname{supp}\mu$. ■

Proof of Proposition 5. By Proposition 4 $\succsim^h$ admits an average robust control representation $\left(u, Q, \mu\left(\cdot|h\right), \lambda_h\right)$ for every $h \in \mathcal{H}$. Observe that since the outcome frequency is constant along the sequence $\left(h_{t_n}\right)_{n\in\mathbb{N}}$, 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 $\rho \in \Delta\left(\Omega\right) \setminus \left\{\rho_q : q \in Q\right\}$ and let $\succsim_{\lim}$ be the preferences whose existence is guaranteed by the Asymptotic Frequentism. It is immediate to see that the average robust control representation $\left(u_{\lim}, Q_{\lim}, \mu_{\lim}, \lambda_{\lim}\right)$ of $\succsim_{\lim}$ must have $u_{\lim} = u$ and $Q_{\lim} = \operatorname{supp}\mu_{\lim} \subseteq \operatorname{supp}\mu = Q$. Suppose by way of contradiction that

$$\begin{aligned} l: &= \liminf_{n\to\infty} c\lambda^{h_{t_n}} = \liminf_{n\to\infty} \lambda_{h_{t_n}} / \left(\frac{LRT\left(h_{t_n},Q\right)}{2t_n}\right) \\ &< \limsup_{n\to\infty} \lambda_{h_{t_n}} / \left(\frac{LRT\left(h_{t_n},Q\right)}{2t_n}\right) = \limsup_{n\to\infty} c\lambda_{h_{t_n}} =: L. \end{aligned}$$

Then either there is an $\varepsilon \in \mathbb{R}_{++}$ with $\left(\lambda_{\lim} - \varepsilon\right)c > l$ or there is an $\varepsilon \in \mathbb{R}_{++}$ with $\left(\lambda_{\lim} + \varepsilon\right)c < L$. We obtain a contradiction in the first case, the proof for the other case being symmetric.

Recall that since the state space is Q-adequate there exist

$$k \in (0,1) \text{ and } \left(W_q\right)_{q\in Q} \in \left(2^{\Omega}\right)^{Q}$$

such that $\rho_q(W_q) = k$. Let $x, y, z \in X$ be such that $x \succ^{\emptyset} z$ and

$$\begin{aligned} & -k\exp\left(-\left(\lambda_{\lim}-\varepsilon\right)\left(u\left(z\right)\right)\right)-\left(1-k\right)\exp\left(-\left(\lambda_{\lim}-\varepsilon\right)\left(u\left(x\right)\right)\right) \\ = & -\exp\left(-\left(\lambda_{\lim}-\varepsilon\right)\left(u\left(y\right)\right)\right), \end{aligned}$$

where the existence of such y is guaranteed by u being affine, and clearly $x \succ^{\emptyset} y$. Define $E = \cup_{q\in Q}\left\{\rho_q\right\} \times W_q$. Consider a subsequence $(n_m)_{m\in\mathbb{N}}$ such that

$$\lim_{m\to\infty} c\lambda^{h_{t_{n_m}}} = l.$$

Moreover, let $M \in \mathbb{N}$ be such that for all $m \geq M$

$$c\lambda^{h_{t_{n_m}}} < \frac{\left(\lambda_{\lim}-\varepsilon\right)c+l}{2}.$$

With this, by Proposition 4 and Lemma 14 we have that for all $m \geq M$

$$zEx \prec_{\lim} y = yEy \text{ and } zEx \succ^{h_{t_{n_m}}} y = yEy.$$

But this in turn implies that $\succsim^{h_{t_{n_m}}}$ is never (x, y, E) similar to $\succsim_{\lim}$ for $m \geq M$, a contradiction. This shows that either $\lambda_{h_{t_n}}$ converges or it diverges to plus infinity.

Suppose by contradiction that $\lambda_{h_{t_n}}$ diverges to plus infinity and (possibly restricting to a subsequence) suppose it is strictly increasing. Let $x, y \in X$ be such that $x \succ^{\emptyset} y$ and such that there exist $\bar{z}, \underline{z} \in X$ with $\bar{z} \succ x$ and $y \succ \underline{z}$. Let $\hat{q} \in Q$ be such that $\Omega \times \left\{\rho_{\hat{q}}\right\}$ is nonnull under $\succsim^{\emptyset}$. Then, it contains at least three disjoint nonnull events $W \times \left\{\rho_{\hat{q}}\right\}$, $W' \times \left\{\rho_{\hat{q}}\right\}$, and $W'' \times \left\{\rho_{\hat{q}}\right\}$, and without loss of generality take them such that they partition $\Omega \times \left\{\rho_{\hat{q}}\right\}$ as well as with $\rho_{\hat{q}}(W) \leq \frac{1}{3}$. We now show that for every $M \in \mathbb{N}$ there is an $n \geq M$ such that $\succsim^{h_{t_n}}$ is not $\left(x, y, W \times \left\{\rho_{\hat{q}}\right\}\right)$ similar to $\succsim_{\lim}$. Let $z, z', z'' \in X$ be such that $y \succ z'$, $z \succ z'$, and

$$\begin{aligned} & -\left(1-\rho_{\hat{q}}\left(W\right)\right)\exp\left(-\lambda_{\lim}\left(u\left(z\right)\right)\right)-\rho_{\hat{q}}\left(W\right)\exp\left(-\lambda_{\lim}\left(u\left(y\right)\right)\right) \\ < & -\rho_{\hat{q}}\left(W'\right)\exp\left(-\lambda_{\lim}\left(u\left(z'\right)\right)\right)-\rho_{\hat{q}}\left(W''\right)\exp\left(-\lambda_{\lim}\left(u\left(z''\right)\right)\right)-\rho_{\hat{q}}\left(W\right)\exp\left(-\lambda_{\lim}\left(u\left(x\right)\right)\right) \end{aligned}$$

where the existence of such z, z', z'' is guaranteed by u being affine. By Proposition 4

and Lemma 14 there exists an $\hat{L} \in \mathbb{R}_{++}$ such that for all $L > \hat{L}$

$$\begin{aligned} & -\left(1-\rho_{\hat{q}}\left(W\right)\right)\exp\left(-L\left(u\left(z\right)\right)\right)-\rho_{\hat{q}}\left(W\right)\exp\left(-L\left(u\left(y\right)\right)\right) \\ > \quad & -\rho_{\hat{q}}\left(W'\right)\exp\left(-L\left(u\left(z'\right)\right)\right)-\rho_{\hat{q}}\left(W''\right)\exp\left(-L\left(u\left(z''\right)\right)\right)-\rho_{\hat{q}}\left(W\right)\exp\left(-L\left(u\left(x\right)\right)\right). \end{aligned}$$

Therefore, we have that

$$y\left(W\times\left\{\rho_{\hat{q}}\right\}\right)z \prec_{\lim} x\left(W\times\left\{\rho_{\hat{q}}\right\}\right)\left(z'\left(W'\times\left\{\rho_{\hat{q}}\right\}\right)z''\right)$$

and

$$y\left(W\times\left\{\rho_{\hat{q}}\right\}\right)z \succ^{h_{t_{n_m}}} x\left(W\times\left\{\rho_{\hat{q}}\right\}\right)\left(z'\left(W'\times\left\{\rho_{\hat{q}}\right\}\right)z''\right)$$

for every h_{t_n} for which $\lambda_{h_{t_n}} > L$. Since by assumption $\lambda_{h_{t_n}}$ diverges to plus infinity, we have obtained a contradiction. ■

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

$$L\colon = \limsup_{n\to\infty}\frac{LRT\left(h_{t_n},Q\right)}{\lambda_{h_{t_n}}t_n} > 0 \tag{30}$$

for some sequence of histories with empirical frequency $\hat{\rho} \notin \left\{\rho_q : q \in Q\right\}$. Then by Structured Savage and equation (29) there exists $\varepsilon > 0$ such that $\Omega \times B_\varepsilon(\hat{\rho})$ is $\succsim^\emptyset$-null. Since the state space is Q-adequate there exist

$$k \in (0,1) \text{ and } (W_q)_{q\in Q} \in \left(2^\Omega\right)^Q$$

such that $\rho_q(W_q) = k$ for all $q \in Q$. Define $E = \cup_{q\in Q}\left(W_q \times \left\{\rho_q\right\}\right)$. 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

$$-\exp\left(-2\frac{c}{L}u\left(z\right)\right) = -k\exp\left(-2\frac{c}{L}u\left(x\right)\right)-\left(1-k\right)\exp\left(-2\frac{c}{L}u\left(y\right)\right)$$

where the existence of such z is guaranteed by u being affine and X being convex. But

then equation (30) implies that for infinitely many $n \in \mathbb{N}$

$$f \succ^{h_{t_n}} z$$

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 for some $K \in \mathbb{R}$

$$L := \liminf_{n\to\infty} \frac{LRT(h_{t_n}, Q)}{\lambda_{h_{t_n}} t_n} < K. \tag{31}$$

As the state space is Q-adequate there exist $k \in (0,1)$ and $(W_q)_{q\in Q} \in \left(2^{\Omega}\right)^Q$ such that $\rho_q(W_q) = k$. Let $x, y, z \in X$ be such that $x \succ^{\emptyset} z$ and

$$\begin{aligned} & -k\exp\left(-\frac{c}{2\max\{1,L\}}(u(z))\right) - (1-k)\exp\left(-\frac{c}{2\max\{1,L\}}(u(x))\right) \\ = & -\exp\left(-\frac{c}{2\max\{1,L\}}(u(y))\right), \end{aligned}$$

where the existence of such y is guaranteed by u being affine, and clearly $x \succ y$. Define $E = \cup_{q\in Q}\left(W_q \times \{\rho_q\}\right)$. Consider a subsequence $(n_m)_{m\in\mathbb{N}}$ such that

$$\lim_{m\to\infty} \frac{LRT\left(h_{t_{n_m}}, Q\right)}{\lambda_{h_{t_{n_m}}} t_{n_m}} = L.$$

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 $m \geq M$

$$zEx \prec^{h_{t_{n_m}}} y = yEy \text{ and } zEx >^{h_{t_{n_m}}} y = yEy.$$

By Corollary 2, this contradicts Asymptotic Leniency. ■

A.3 General Statistical Distances

The results of the paper that involve the statistician type extend easily also 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

$$\mathbb{E}_{\mu}\left[\min_{p\in\Delta(S)}\mathbb{E}_{p}\left[u\left(f\right)\right]+\frac{1}{\lambda}D_{\phi}\left(p||q\right)\right]$$

where

$$D_{\phi}\left(p||q\right)=\begin{cases}\int\phi\left(\frac{\mathrm{d}p}{\mathrm{d}q}\right)\mathrm{d}q & q\gg p\\ \infty & \text{otherwise}\end{cases}$$

for some continuous strictly convex function $\phi:\mathbb{R}_{+}\to\mathbb{R}_{+}$ with

$$\phi\left(1\right)=0\text{ and }\lim_{t\to\infty}\frac{\phi\left(t\right)}{t}=\infty.$$

From this expression, it is clear that the main case studied in the paper is the one where $\phi\left(c\right)=c\log c-c+1$. 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 Example 1

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

$$\mathbb{E}_{\omega,\varepsilon_2}\left[\left(a+\omega\right)\left(1-\theta-\varepsilon_2\right)-\left(a'+\omega\right)\left(1-\theta-\varepsilon_2\right)\right]\geq c\left(a\right)-c\left(a'\right).$$

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

$$-\frac{k}{R\left(p_a^*||q_a^{\theta}\right)}\log\mathbb{E}_{\omega,\varepsilon_2}\left[\exp\left(-\frac{k}{R\left(p_a^*||q_a^{\theta}\right)}\left[\left(a+\omega\right)\left(1-\theta-\varepsilon_2\right)-\left(a'+\omega\right)\left(1-\theta-\varepsilon_2\right)\right]\right)\right]\geq c\left(a\right)-c\left(a'\right).$$

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 Berlin and Heidelberg GmbH & Company KG.

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*.

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.

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.

Bohren, J. A. and D. N. Hauser (2021). "Learning with heterogeneous misspecified models: Characterization and robustness". *Econometrica* 89, pp. 3025–3077.

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.

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.

Denti, T. and L. Pomatto (2022). "Model and predictive uncertainty: A foundation for smooth ambiguity preferences". *Econometrica* 90, pp. 551–584.

Diaconis, P. and D. Freedman (1986). "On the Consistency of Bayes Estimates". *The Annals of Statistics* 14, pp. 1–26.

Diamond, P. A. (1982). "Aggregate demand management in search equilibrium". *Journal of political Economy* 90, pp. 881–894.

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.

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 (2021). "Asymptotic behavior of Bayesian learners with misspecified models". *Journal of Economic Theory* 195, p. 105260.

Esponda, I., E. Vespa, and S. Yuksel (2020). *Mental models and learning: The case of base-rate neglect*. Tech. rep. Tech. rep.

Fishburn, P. C. (1970). *Utility theory for decision making*. Wiley.

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.

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?" *Available at SSRN 3709932*.

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: Journal of the Econometric Society*, 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.

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.

— (2022). *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.

He, K. and J. Libgober (2022). "Evolutionarily stable (mis) specifications: Theory and applications".

Herstein, I. N. and J. Milnor (1953). "An axiomatic approach to measurable utility". *Econometrica, Journal of the Econometric Society*, pp. 291–297.

Kallenberg, O. (1973). "Random measures". *Random Measures*. De Gruyter.

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.

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.

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.

Roberts, A. and D. Varberg (1973). *Convex Functions*. Academic Press.

Royden, H. L. and P. Fitzpatrick (1988). *Real analysis*. Vol. 32. Macmillan New York.

Sargent, T. J. (1999). "The conquest of American inflation". *The Conquest of American Inflation*. Princeton University Press.

— (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.

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: Journal of the Econometric Society*, pp. 307–333.

Wakker, P. P. (2013). *Additive representations of preferences: A new foundation of decision analysis*. Vol. 4. Springer Science & Business Media.