

Dynamic Concern for Misspecification*

Giacomo Lanzani

Harvard

Abstract

I 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. I show that different static preferences under uncertainty (subjective expected utility, maxmin, robust control) arise in the long run, depending on how quickly the agent becomes unsatisfied with unexplained evidence. The misspecification concern's endogeneity naturally induces behavior cycles, and I characterize the limit action frequency. I apply the model to monetary policy cycles and choices in the face of complex tax schedules.

*I am particularly indebted to Simone Cerreia-Vioglio, Roberto Corrao, Drew Fudenberg, Stephen Morris, and Philipp Strack. I am also grateful to Pierpaolo Battigalli, Stéphane Bonhomme, Benjamin Brooks, Yeon-Koo Che, Daniel Clark, Tommaso Denti, David Dillenberger, Federico Echenique, Andrew Ellis, Glenn Ellison, Joel Flynn, Mira Frick, Benjamin Golub, Olivier Gossner, Faruk Gul, Lars Peter Hansen, Johannes Hörner, Ryota Iijima, Peter Klibanoff, Asen Kochov, Jonathan Libgober, Fabio Maccheroni, Massimo Marinacci, Eric Maskin, Moritz Meyerter-Vehn, Anh Nguyen, Jawwad Noor, Parag Pathak, Luciano Pomatto, Damien Pouzo, Giulio Principi, Phil Reny, Vasiliki Skreta, Karl Schlag, Advik Shreekumar, Marciano Siniscalchi, Tomasz Strzalecki, Jean Tirole, Anton Tsoy, Emanuel Vespa, Alexander Wolitzky, Yuichi Yamamoto, and Leeat Yariv 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 (DGP) postulates multiple probabilistic descriptions of the environment and uses Bayes rule to adjust their relative weights. However, even rational agents may fear that they are misspecified and that none of these descriptions is correct. This concern is remarkably natural in complex and high-dimensional settings, where uncertainty needs to be simplified to obtain well-behaved optimization and learning procedures.

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

All these situations share the use of structured models expressed in terms of interpretable parameters, such as the slope of the Phillips curve, demand elasticities, or the descent rate of raindrops. The best among these parametrized models can be generally learned, but the agent may be concerned with misspecification, i.e., with all these models being too simple to be correct.

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 et al., 2021, and the references therein). These works assume that the agents have no concern about being misspecified. Here, I 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 uncertainty-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), whose decision criterion has been axiomatized

by Strzalecki (2011), considers agents who fear model misspecification.¹ In particular, the first axiomatization that separates objective risk, uncertainty aversion, and concern for misspecification has been proposed in Cerreia-Vioglio et al. (2024). This work reconciles these approaches and shows how popular decision criteria such as maxmin expected utility, robust control preferences, and SEU arise as the limit behavior of an agent concerned about misspecification and learning about the actual DGP.

I consider an agent who repeatedly chooses among actions whose payoffs have an unknown distribution. This choice is taken by evaluating each action with respect to an average of robust control assessments, where each assessment takes a different structured model as the benchmark. I introduce endogeneity in the misspecification concern: the better the structured models explain the past, the less concerned the agent is.

To characterize the limit behavior under misspecification, a taxonomy of how demanding the agent is in evaluating the explanatory performance of their model turns out to be crucial. I develop this taxonomy by first establishing a normative benchmark: a “statistically sophisticated” adjustment of the concern for misspecification that performs a likelihood-ratio evaluation of the agent’s structured models. More precisely, I show that achieving two desirable properties uniquely characterizes such behavior: global safety (i.e., guaranteeing the minmax payoff even when misspecified) and self-confirming under correct specification (i.e., no loss beyond those caused by the classical self-confirming learning traps when the agent is correctly specified, cf. Easley and Kiefer, 1988).

I then consider departures from this normative benchmark to obtain descriptive predictions on the effect of an endogenous misspecification concern. I consider agents that are too demanding in evaluating the structured models’ performance (this case includes believers in the Law of Small Numbers, Tversky and Kahneman, 1971, that treat failures in explaining early realizations as a statistician treats long-run failures). Similarly, I allow the opposite case in which the agent is too lenient in evaluating their structured models and attributes too much unexplained evidence to sampling variability.

Specifically, I characterize the long-run behavior of these different types of misspecified agents. The actions of the lenient type converge to a Berk-Nash equilibrium, i.e., to an SEU best reply to beliefs supported on the structured models closest to the actual DGP. Instead, overemphasis on

¹Strzalecki (2011) can be additionally interpreted as providing an axiomatization of robust prior analysis, see Hansen and Sargent (2022a).

the structured models' failures in explaining the data by the demanding type induces convergence to a maxmin best reply to the DGPs that are absolutely continuous with respect to the true one.

In contrast, a statistically sophisticated type maintains a nonextreme concern for misspecification. If their behavior converges, it converges to a robust control best reply to the structured models closest to the actual DGP. Moreover, the misspecification concern is endogenously determined by how well the best structured 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.² First, the extent of long-run uncertainty aversion positively correlates with the agent being initially misspecified and their belief in the Law of Small Numbers. 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, the limit uncertainty attitudes are stochastic, even keeping the misspecification and understanding of probability rules constant. Initial realizations leading to a limit action with consequences poorly explained by the agent's structured models induce a long-run uncertainty aversion higher than realizations, leading to a limit action whose consequences are well explained.

I then investigate the impact of an endogenous misspecification concern in a class of problems where a complex, nonlinear function describing the average returns from a costly action is simplified to a linear one. I show that if the actual average returns are concave, an endogenous misspecification concern helps. This result covers incorrect beliefs in price taking (Sobel, 1984) and simplification of complex tax schedules (Esponda and Pouzo, 2016). For the latter application, I rationalize the labor supply in the face of complex tax schedules documented in Rees-Jones and Taubinsky (2020). In particular, they show that around 40% of the agents have beliefs corresponding to a heuristic that simplifies the tax schedule to a linear one but that 20% fewer agents act according to this heuristic. This is predicted by an endogenous concern for misspecification, as agents with an incorrect structured model are less prone to base their decisions on the conclusions they reach within the model.

In general, the behavior of a statistically sophisticated type is not guaranteed to converge. Indeed, it is possible that sophisticated types behavior cycles between phases of different misspecification concerns. Still, I characterize the limit action frequency and concern for misspecification. I apply this result to revisit the cyclical behavior of monetary policies documented in Sargent

²The empirical study of the correlation between behavioral biases is an active area of recent development. See, e.g., Dean and Ortleva (2019), Snowberg and Yariv (2021), and the references therein.

(2008). Intuitively, the cycles have the following structure. The agent spends some time playing an action whose consequences are well explained by one of their structured models (a conservative monetary policy in the application). Playing this action lowers the concern for misspecification and eventually leads to a more misspecification-vulnerable action (a more aggressive monetary policy). Failures to explain the distribution of outcomes observed under this action lead to a return to the more misspecification robust action.

1.1 Related Literature

Paradigm Change A few papers allow the agents to realize that they are misspecified. In particular, in Cho and Libgober (2023), Fudenberg and Lanzani (2023), Gagnon-Bartsch et al. (2023), Ba (2024), and He and Libgober (2024), misspecification can be eliminated by “light bulb realizations” evolutionary pressure, or the use of a machine-learning-inspired algorithm. The critical difference with our approach is that in these papers, as well as in the earlier Cho and Kasa (2015) and Giacomini et al. (2020), where agents switch between models based on a specification test, the agents act as if they have complete trust in the set of models currently entertained and are never concerned about being misspecified.

Still, there is a tight connection between the robust control decision criterion and a maxmin decision criterion where the set of models expands as the penalization term in the robust control increases (see Hansen and Sargent, 2011 for a textbook treatment). In light of this, compared to the previous set of papers, our work can additionally be interpreted as providing a *smooth* framework for expanding (or restricting) the set of considered models as a function of the evidence.

Decision Theory The static decision criterion considered here is due to Cerreia-Vioglio et al. (2024). The explicit use of a state space where every state describes both the single-period outcome realization and the probability distribution over outcomes follows the approach introduced in Cerreia-Vioglio et al. (2013) as a two-stage “statistical” interpretation and axiomatization of some of the decision criteria under ambiguity, in particular the smooth ambiguity one.

Although not dealing with misspecification, Epstein and Schneider (2007) also study a learning problem with nonSEU preferences and where the likelihood ratio between different DGPs plays an important role. I discuss the difference between the two approaches in Remark 2. Dillenberger and Rozen (2015) propose a model where risk attitudes are influenced by the performance of the decision maker’s past actions and show how it induces higher volatility and path dependence in a

dynamic asset pricing model.

Misspecified Learning There is fast-growing literature on learning under misspecification with SEU preferences. Arrow and Green (1973) give the first general framework for this problem, and Nyarko (1991) points out that the combination of misspecification and endogenous data can lead to cycles. This literature has been revived by the more recent Esponda and Pouzo (2016); see Bohren and Hauser (2021), Esponda et al. (2021a), Fudenberg et al. (2021), Frick et al. (2023), and Fudenberg et al. (2024) for analyses of the most closely related settings. The characterization of the limit beliefs borrows extensively from their insights.

Determining the Concern for Misspecification Hansen and Sargent (2007) consider 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. Anderson et al. (2003) and Barillas et al. (2009) pioneer a literature that calibrates the (time-invariant) concern for misspecification from the acceptable error probability in a likelihood-ratio test between the unperturbed model and the worst-case model (that does not depend on the action there). See Hansen and Sargent (2011) for a textbook treatment.

2 Decision Criterion

2.1 Static Decision Criterion

I consider an agent who chooses from a finite number of actions $a \in A$ and let Y be a compact metric space representing the set of possible outcomes. The agent has a continuous utility index $u : A \times Y \rightarrow \mathbb{R}$ over the action-outcome pairs that captures their preference when the 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.³

³Except for Section 2 of the Supplemental Material, where another topology over probability measures is also considered, for every subset C of a metric space, I denote as $\Delta(C)$ the Borel probability measures on C , endowed with the topology of weak convergence of measures.

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 $Q \subseteq \Delta(Y)^A$ of *structured models*, i.e., action-dependent probability measures over outcomes $q = (q_a)_{a \in A}$. The agent has a prior belief $\mu \in \Delta(Q)$ with support Q that describes the relative likelihood assigned to these models.

The interpretation of these structured models is that they are the simplified models that the DM uses to make predictions and infer from what they observe. In our applications, 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 without a systematic effect of inflation on unemployment. Alternatively they may be a taxpayer deciding how much to work using a simplified and reduced version of the complex tax code they face. More generally, structured models differ from arbitrary probability distributions over states by being formulated in terms of parameters that relate actions and outcomes with a concrete economic interpretation. Often, this makes them expressible as a finite dimensional set of models, but the regularity conditions, I assume, do not, in general, impose finite dimensionality.⁴

Assumption 1. For every $a \in A$: (i) For all $q \in Q$, $\tilde{q}_a := \frac{dq_a}{dp_a^*}$ exists, is continuous, and is p_a^* -a.s. bounded away from 0, uniformly in Q , (ii) For p_a^* -almost every $y \in Y$ the map $q \mapsto \tilde{q}_a(y)$ is continuous.

Condition (i) allows us to compute the relevant expectations while allowing for discrete and continuous outcome spaces. It guarantees that no subjective model of the agent is ruled out in finite time.⁵ Continuity of the map from models to outcome distributions is a standard requirement for parametric models.

A myopic Bayesian agent with complete trust in their models evaluates action a according to its SEU:

$$\int_Q \mathbb{E}_{q_a} [u(a, y)] d\mu(q).$$

⁴The terminology of structured models first appeared in Hansen and Sargent (2022b), where, in their words, they chose to “call models “structured” because they are parsimoniously parameterized based on a priori considerations.”

⁵Part (i) also plays a technical role in two proofs. It helps guarantee the existence of the equilibrium concepts I consider (although it is known it can be relaxed; see Anderson et al., 2022). Moreover, it allows us to use the techniques developed for establishing convergence with random measures (see, e.g., Kallenberg, 1973). Again, that literature can deal with weaker but less explicit conditions.

That is, they compute a two-stage expectation of the utility function: they evaluate the utility of the action given the candidate model q , $\mathbb{E}_{q_a} [u(a, y)]$, and then they average over the models with weights given by their subjective belief μ .⁶

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

The agent evaluates each action a according to the *average robust control* criterion:

$$\int_Q \min_{p_a \in \Delta(Y)} \left(\mathbb{E}_{p_a} [u(a, y)] + \frac{1}{\lambda} R(p_a || q_a) \right) d\mu(q) \quad (1)$$

where $\lambda > 0$ is a parameter that trades off between decision robustness and performance under the structured models.

The original robust control preference 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 true DGP p^*). As described in Hansen et al. (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) follows Hansen and Sargent (2007) and Cerreia-Vioglio et al. (2024) in extending this interpretation to a situation in which the agent is still uncertain about the best-approximating structured model (i.e., μ is nondegenerate), allowing the malevolent nature to alter each of the candidate structured models.

The representation adopts the distinction between two levels of uncertainty. At the first level, given a structured model q , the uncertainty about the exact specification of the model is captured by minimizing the expected utility for probabilities that are not too far away from q . 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 SEU, the lower level captures the agent’s concern for misspecification. Observe that by considering a *simple*

⁶The myopic assumption allows us to underscore our main points without dealing with issues of dynamic consistency. See Section 5.1 for a detailed discussion of the extension to forward-looking agents.

average of robust control evaluations, rather than a concave transformation ϕ of those values, I focus on a decision maker who is concerned about misspecification but it is ambiguity neutral (see also the detailed discussion in Cerreia-Vioglio et al. (2024) and Hansen and Sargent (2022a)).⁷

Remark 1. [Axiomatization] A full-fledged axiomatization of this decision criterion is obtained in Section 2 of the Supplemental Material. This remark summarizes the key axioms. In terms of observability requirements, I allow the analyst to elicit preferences for bets on the DGP, e.g., the urn composition, and on the actual realization, e.g., the color of the drawn ball.⁸

The static decision criterion belongs to the variational class of Maccheroni et al. (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 the urn composition, see Axiom 2) and for bets on events given the structured model (e.g., bets on the color after having revealed the urn composition, see Axiom 3, here the conclusion follows directly from Strzalecki (2011)). However, failures of the Sure-Thing principle are allowed for acts that involve the realization of the outcome without conditioning on the model (e.g., bets on the color without knowing the urn composition, which are the ones involved in the classical Ellsberg’s paradox). The final conceptual axiom involved in the representation of equation (1) is a notion of uniform conditional misspecification concern (see Axiom 4). It requires that conditional on being told the identity of their best-fitting structured model, the agent is equally concerned about it not being exact regardless of which one it is. I discuss extensively this axiom and modelling decision in Section 5.2. ▲

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, the behavior responds to the received information. Formally, time is discrete, and a history is a finite vector of past actions and outcomes. In particular, the set of histories of length $t \in \mathbb{N}$ is $\mathcal{H}_t = (A \times Y)^t$, and the set of all histories is $\mathcal{H} = \bigcup_{t=0}^{\infty} \mathcal{H}_t$. I will denote with $\mathbf{a}_t, \mathbf{y}_t$, and \mathbf{h}_t the random variables corresponding to the action, outcome, and history at time t , and I use the non-bold version for their realizations.

On the one hand, I stick to the classical treatment of beliefs about the possible DGPs. I let

⁷See Battigalli et al. (2019) for a thorough analysis of ambiguity aversion in the same repeated problem setting of this paper. The combined study of ambiguity aversion and concern for misspecification is left for future work.

⁸This is standard when dealing with multiple sources of uncertainty, see for example Klibanoff et al. (2005), and Cerreia-Vioglio et al. (2013) for a general framework and results.

the belief be updated through standard Bayesian updating. That is, for every measurable subset C of Q , I denote by

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

the subjective belief the agent obtains using Bayes rule after history $(a^t, y^t) \in \mathcal{H}_t$.⁹

On the other hand, I introduce an endogenous and time-evolving concern for misspecification, i.e., I let λ depend on the realized history through a function $\Lambda : \mathcal{H} \rightarrow \mathbb{R}_+$.

Log-Likelihood Ratio In particular, I want to capture the idea that the concern for misspecification is a function of how well the structured models explain the current history. In statistics, a standard measure of fit of a set of distributions Q against a set of unstructured alternatives $N(Q) \subseteq \Delta(Y)^A$ is the log-likelihood ratio:

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

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

Assumption 2. (i) $p^* \in N(Q) \supseteq Q$. (ii) For every $a \in A$, the family of densities $\{\tilde{p}_a : p \in N(Q)\}$ is a compact set of continuous functions.

I require that the unstructured set is a relaxation sufficiently large to include the actual distribution and a continuity condition that rules out a Q that only contains continuous distributions

⁹By Assumption 1 (i), the posterior is well-defined except on a set of histories with 0 objective probability. I allow for arbitrary belief revisions after those histories.

and an $N(Q)$ that includes discrete distributions.¹⁰

An important role is played by the rule

$$\Lambda(h_t) = \frac{\text{LLR}(h_t, Q)}{ct} \quad \forall t \in \mathbb{N}, \forall h_t \in \mathcal{H}_t \quad (2)$$

where $c > 0$. Under this rule, the concern for misspecification is proportional to the average log-likelihood ratio. This *average* log-likelihood ratio is proposed in statistics to measure the extent of model misspecification. Therefore, I refer to an agent who uses such a rule as a “statistically sophisticated type”. Of course, the estimation goal of a statistician can be very different from that of an agent involved in a repeated decision problem under uncertainty. Theorem 1 below establishes that this rule is also a rationality benchmark in repeated decision problems.

3 Long-run Payoffs and Actions

This section studies the long-run consequences of using the average robust control decision criterion. Our primary interest is in what attitudes towards unexplained evidence, i.e., what Λ , induce good payoff performance across environments and what are the limit actions and preferences under uncertainty that arise given a specific attitude.

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

$$BR^\lambda(\nu) = \operatorname{argmax}_{a \in A} \int_Q \min_{p_a \in \Delta(Y)} \left(\mathbb{E}_{p_a} [u(a, y)] + \frac{R(p_a || q_a)}{\lambda} \right) d\nu(q) \quad \forall \lambda \in \mathbb{R}_{++}, \forall \nu \in \Delta(Q).$$

Also, let

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

¹⁰Assumption 2 (ii) generally requires that the DM imposes some restrictions on the unstructured models, i.e., $N(Q) \neq \Delta(Y)^A$ while Assumption 2 (i) requires the restriction to be sufficiently permissive to include p^* . This creates a tension that is especially undesirable for the normative Theorem 1. A reasonable preoccupation is whether the LLR-based adjustment of the concern for misspecification performs well only because the agent is assumed to be able to guess a correctly specified set of unstructured alternatives. In the Working Paper version, I show that the answer turns out to be negative, as Theorem 1 continues to hold when (i) is suitably relaxed.

denote the actions that maximize the SEU of an agent with belief ν and

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

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

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

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

3.1 Normative Benchmark

Now, I provide a normative justification for considering the log-likelihood ratio rule (2) the relevant benchmark of rationality, showing that it satisfies two desirable properties across the possible decision problems.

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

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

This is a very mild condition that only requires the agent to obtain an average payoff at least ε close to what they can guarantee against *every* possible outcome. However, when paired with misspecification, ε -safety has a significant bite: a Bayesian SEU agent fails it in many decision problems, as illustrated by Example 5 in the Supplemental Material. Indeed, such failures have been the basis of many critiques of the predictions of learning under misspecification with Bayesian SEU agents.

The next property requires that, in the face of sufficiently small misspecification, Λ does not induce any performance loss compared to the possible self-confirming learning traps that are well-known (Battigalli, 1987, Easley and Kiefer, 1988, and Fudenberg and Levine, 1993) to affect even correctly specified SEU agents who are not infinitely patient. Formally, for every $\varepsilon > 0$, define the

ε -self-confirming correspondence as

$$SCE_\varepsilon(p^*) = \left\{ a \in A : \exists \nu \in \Delta(\{q \in Q : R(p_a^*||q_a) \leq \varepsilon\}), \int_Q \mathbb{E}_{q_a} [u(a, y)] d\nu(q) + \varepsilon \geq \max_{a' \in A} \int_Q \mathbb{E}_{q_{a'}} [u(a', y)] d\nu(q) \right\}$$

and the ε -self-confirming guarantee as $G_\varepsilon(p^*) = \min_{a \in SCE_\varepsilon(p^*)} \mathbb{E}_{p_a^*} [u(a, y)]$.

In other words, $SCE_\varepsilon(p^*)$ contains all the ε -best replies to beliefs that come ε -close to be confirmed on path, while $G_\varepsilon(p^*)$ is the minimal objective performance among the actions in $SCE_\varepsilon(p^*)$.

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

$$\min_{q \in Q} \max_{a \in A} R(p_a^*||q_a) < \delta \implies \liminf_{t \rightarrow \infty} \frac{\sum_{i=1}^t u(\mathbf{a}_i, \mathbf{y}_i)}{t} \geq G_\varepsilon(p^*) - \varepsilon \quad \mathbb{P}_\Pi\text{-a.s.}$$

ε -self-confirming under correct specification requires that sufficiently low levels of misspecification (i.e., the existence of a model q with distance less than δ from the actual data-generating process) induce arbitrarily small losses beyond those that could be already incurred by a correctly specified SEU agent (i.e., a limit average payoff more than ε lower than the expected payoff of a self-confirming equilibrium). Intuitively, if the misspecification is minor, then the agent approximately identifies the on-path consequences of the true DGP and starts best replying to them in the long run. In particular, whenever the environment is such that there exists a unique best explaining model independent of the action, ε -self-confirming under correct specification means that the DM achieves no regret in the long run, with an average payoff equal to the ex-post optimal one.¹¹

Theorem 1. 1. For every $\varepsilon > 0$ and decision problem there exists $c > 0$ such that if

$$\Lambda(h_t) = \frac{LLR(h_t, Q)}{ct} \quad \forall t \in \mathbb{N}, \forall h_t \in \mathcal{H}_t,$$

then Λ is both ε -safe and ε -self-confirming under correct specification.

2. There exist $\varepsilon > 0$ and a decision problem for which there is no ε -safe and ε -self-confirming

¹¹See the Working Paper for the formal statement and proof of this special case.

under correct specification Λ with either¹²

$$\lim_{t \rightarrow \infty} \frac{\Lambda(h_t)}{LLR(h_t, Q)/t} = 0 \quad \forall (h_t)_{t \in \mathbb{N}} \in \times_{t \in \mathbb{N}} \mathcal{H}_t,$$

or

$$\lim_{t \rightarrow \infty} \frac{LLR(h_t, Q)/t}{\Lambda(h_t)} = 0 \quad \forall (h_t)_{t \in \mathbb{N}} \in \times_{t \in \mathbb{N}} \mathcal{H}_t.$$

The safety and self-confirming conditions single out the LLR-based adjustment of the misspecification concern as a normatively appealing rule for a statistically sophisticated agent. At the same time, some less normatively appealing but descriptively relevant phenomena are captured by other rules. On the one hand, a rule such that $\lim_{t \rightarrow \infty} \frac{\Lambda(h_t)}{LLR(h_t, Q)/t} = \infty$, e.g., $\Lambda(h_t) = \frac{LLR(h_t, Q)}{\sqrt{t}}$, overly penalizes minor imperfections of the model, expecting that the frequency quickly converges to its theoretical value, as in the Law of Small Numbers fallacy. On the other hand, an agent for which $\lim_{t \rightarrow \infty} \frac{\Lambda(h_t)}{LLR(h_t, Q)/t} = 0$ attributes too much of the unexplained evidence to sampling variability. In this regard, Theorem 1 tells us that there are decision problems where any way to adjust the concern for misspecification that is globally more demanding or lenient than the average LLR violates either safety or self-confirming under correct specification. Moreover, SEU maximization is not safe, while always using a maxmin best reply to Q induces a behavior that is not self-confirming under correct specification.

3.2 Positive Long-run Predictions

In this section, I show that the limit behavior can be described through fixed point conditions involving the agent's action, belief, and concern for misspecification. To this end, let $Q(a) = \operatorname{argmin}_{q \in Q} R(p_a^* || q_a)$ be the structured models that best fit the actual DGP when action a is played.

Definition 4. Action a^* is a:

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

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

¹²I use the convention that $\frac{0}{0} = 0$.

2. *Maxmin equilibrium* if¹³

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

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

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

B-NE (Esponda and Pouzo, 2016) describes a stable situation where the agent's action best replies to a belief that is 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 under every structured model).

c-robust equilibrium is similar to B-NE 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 structured models fit the true DGP at the equilibrium.

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

The following result shows that the taxonomy suggested by Theorem 1 in terms of how quickly the agent becomes unsatisfied with their model is the critical determinant of the long-run agent's behavior.

Theorem 2. Let a^* be a Λ -limit action with $p_{a^*}^* \notin \{q_{a^*}\}_{q \in Q}$.

1. If

$$\lim_{t \rightarrow \infty} \frac{\Lambda(h_t)}{LLR(h_t, Q) / t} = 0 \quad \forall (h_t)_{t \in \mathbb{N}} \in \times_{t \in \mathbb{N}} \mathcal{H}_t, \quad (4)$$

then a^* is a B-NE.

2. If

$$\lim_{t \rightarrow \infty} \frac{LLR(h_t, Q) / t}{\Lambda(h_t)} = 0 \quad \forall (h_t)_{t \in \mathbb{N}} \in \times_{t \in \mathbb{N}} \mathcal{H}_t, \quad (5)$$

then a^* is a maxmin equilibrium.

¹³For every $p, q \in \Delta(Y)$, $p \gg q$ means that q is absolutely continuous with respect to p , and $p \sim q$ means that they are mutually absolutely continuous.

3. If

$$\Lambda(h_t) = \frac{LLR(h_t, Q)}{ct} \quad \forall t \in \mathbb{N}, \forall h_t \in \mathcal{H}_t,$$

then a^* is a c -robust equilibrium.

The theorem characterizes the possible limit actions of all types of agents. At one extreme, the workhorse equilibrium concept of the literature on misspecified learning, B-NE, is sufficient to describe the long-run behavior of *lenient* types. At the other extreme, the repeated failures in explaining the observed data lead demanding agents to a highly pessimistic behavior and consider the worst-case scenario among all the DGPs minimally consistent with the structured models.

Finally, if the behavior of the statistically sophisticated type converges, the limit action best replies 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.

Remark 2 (Dynamic Axiomatization and Epstein and Schneider (2007)). In Section 2 of the Supplemental Material, I consider a collection of binary relations indexed by the observed history to characterize the agent's dynamic preferences. On top of the static axioms discussed in Remark 1, a Dynamic Consistency over Models axiom (see Axiom 7) guarantees that the probability distribution over models is updated in a Bayesian fashion.

Finally, I axiomatize the asymptotic speed of adjustment of the misspecification concern. To do so, I define a quantitative notion of how similar two preference relations are. I use an event E and two deterministic and strictly ranked outcomes, x and y , as measuring rods. Loosely speaking, two relations are (x, y, E, ε) similar if their certain equivalents for the binary act that pays x if E realizes and y otherwise are ε close. With this, an Asymptotic Frequentism axiom (see Axiom 8) singles out the statistically sophisticated type: for every (x, y, E, ε) , the conditional preferences after sufficiently long sequences of outcomes sharing the same empirical frequency must be (x, y, E, ε) -similar.

Although not dealing with misspecification, Epstein and Schneider (2007) also study a learning problem with nonSEU preferences and where the likelihood ratio between different DGPs plays an important role. In the myopic case, and under the notation of our paper, they consider a maxmin decision criterion

$$\min_{\mu \in C_t} \min_{q \in L} \int_{\Theta} \mathbb{E}_{q^\theta} [u(a, y)] d\mu(\theta|h_t)$$

where Θ is a set of parameters, $C_t \subseteq \Delta(\Theta)$, and $L \subseteq \Delta(Y)^\Theta$. Here, the set of possible beliefs over the parameters is recursively given by

$$C_t = \left\{ \nu(\cdot | h_t, q_1, \dots, q_t) : \forall i \in \{1, \dots, t\}, q_i \in L, \nu \in C_0, \frac{\int_{\Theta} \prod_{i=1}^t \tilde{q}_i^\theta(y_i) d\nu(\theta)}{\max_{\mu \in C_0, (p_1, \dots, p_t) \in L^t} \int_{\Theta} \prod_{i=1}^t \tilde{p}_i^\theta(y_i) d\mu(\theta)} \geq \beta \right\},$$

where $\beta \in (0, 1)$ and

$$\nu(B | h_t, q_1, \dots, q_t) = \frac{\int_B \prod_{i=1}^t \tilde{q}_i^\theta(y_i) d\nu(\theta)}{\int_{\Theta} \prod_{i=1}^t \tilde{q}_i^\theta(y_i) d\nu(\theta)}$$

for every measurable B . That is, they look at maxmin preferences with a set of posteriors that only contains the updates from the combination of priors $\nu \in C_0$ and likelihoods $(q_1, \dots, q_t) \in L^t$ that did not perform excessively worse than the best rationalizing one. Crucially, by always considering models that are only supported on Θ , Epstein and Schneider (2007) cannot capture the concern for misspecification that is the core motivation for this paper. It instead captures a gradually increasing confidence in understanding which priors over Θ are the most accurate. This different focus is apparent, inter alia, from the fact that no specification of Epstein and Schneider (2007) satisfies the properties of safety and self-confirming under correct specification.

The papers also have a completely different set of results, as Epstein and Schneider (2007) only deal with the question of beliefs stabilization. Importantly, they are more general in not restricting the agent to be myopic and more specific in assuming that the DGP is exogenous and unaffected by the action chosen by the agent. ▲

3.3 Equilibrium Illustrations

In this section, I revisit some of the main biases that have been justified due to misspecified learning. I first give a result on the positive impact of an endogenous misspecification concern under misspecified linear models. This result is applied to misperceptions of the tax schedule and the competition environment. I then show that some biases can be enhanced, and that is, in particular, the case for correlation neglect.

3.3.1 Misspecified Linear Models

Suppose that $u(a, y) = K_a a - C(a)$, for all $a \in A \subseteq \mathbb{R}_+$, where C is a convex cost function. Here, the misspecification comes from the fact that the agent believes in a simple linear rewards

model with coefficient $K_a = \theta + \varepsilon_{2,a}$, while in reality, $K_a = F(a) + \varepsilon_{1,a}$, where F can be a complicated function, resulting in nonlinear returns from increasing the action. The stochastic terms $\varepsilon_{1,a} \sim N(0, \sigma_1^2)$, $\varepsilon_{2,a} \sim N(0, \sigma_2^2)$ measure respectively actual and conjectured uncertain aspects of the returns to action, and the $(\varepsilon_{1,a}, \varepsilon_{2,a})_{a \in A}$ are independent.¹⁴

Proposition 1. *Suppose that F is a strictly decreasing function. Then, in every Berk-Nash equilibrium a , the action is higher than the optimal one. In every c -equilibrium a' , the bias is reduced, i.e., $a' \leq a$.*

The agent simplifies the (stochastic) return function to a linear form, neglecting the diminishing returns from effort. The estimated constant return is the average return up to the equilibrium action, that is lower than the marginal return. Thus, by equating this perceived marginal return to the marginal cost of effort, the agent ends up playing an excessively higher action. The endogenous misspecification concern reduces the bias because, in the face of unexplained evidence, the agent wants to reduce their exposure to a stochastic process they realize they do not understand fully. Since by playing a lower action, the uncertainty in the returns gets multiplied by a smaller coefficient, this exposure reduction is achieved by decreasing the action.

The assumption of decreasing returns is somewhat abstract. However, I next show how natural forms of misspecification studied in the literature fall in this class and explain how Proposition 1 translates in those settings. The first example shows how the endogenous misspecification concern moderates the Berk-Nash equilibrium's prediction that a more complicated tax schedule induces a higher labor supply.

Example 1 (Bias Reduction under Misperceived Taxation, Esponda and Pouzo, 2016). *An agent chooses a target gross income $a \in A$ at cost $C(a)$. The agent pays taxes at an average tax rate $\tau(a) + \varepsilon_{1,a}$, where $\tau : \mathbb{R} \rightarrow \mathbb{R}$ is strictly increasing as a consequence of a progressive tax schedule. Here $y = t = (\tau(a) + \varepsilon_{1,a})a$, and the payoff is $u(a, y) = a - t - C(a)$. The agent believes in a random coefficient model, $t = (\theta + \varepsilon_{2,a})a$, in which the marginal and average tax rates are equal. See Liebman and Zeckhauser (2004) and Rees-Jones and Taubinsky (2020) for the empirical evidence supporting this “schmeduling” bias.*

¹⁴Formally, ε normally distributed implies that Y is not compact, in contrast with the primary analysis of the paper. Still, the conclusions below are unaffected by considering ε with a symmetrically truncated normal distribution that allows remaining in our main framework. The more natural case is the one where $\sigma_2^2 > \sigma_1^2$, with the interpretation that the agent acknowledges the simplification of their model by allowing a larger volatility. However, Proposition 1 holds for every combination of parameters, and a positive discrepancy between the conjectured and observed behavior (and therefore a possibly strict reduction of the bias in Proposition 1) realizes whenever $\sigma_2^2 \neq \sigma_1^2$.

Here, $Q(a) = \{q^{\tau(a)}\}$, i.e., the best-fitting marginal taxation is equal to the (lower) average taxation. Therefore, in any Berk-Nash equilibrium, the agent exerts higher effort than the optimal.

By Proposition 1, in every c -robust equilibrium, this bias is reduced. To see this, observe that when the agent is not perfectly able to explain the equilibrium data, i.e., $\min_{\theta \in \Theta} R(p_a^* || q_a^\theta) > 0$, they maintain a positive level of concern for misspecification. However, higher efforts are perceived as more exposed to the uncertainty in the marginal rate (as the stochastic tax rate gets multiplied by a higher a).

Therefore, c -robust equilibrium provides a natural force that reduces the counterintuitive prediction that complicated nonlinear taxation codes induce more effort: failures to rationalize the received tax bill reduce effort.¹⁵ This set of predictions is consistent with Rees-Jones and Taubinsky (2020), where it is shown that around 40% of the agents have beliefs (elicited in an incentive-compatible way) corresponding with the scheduling heuristic but that there are 20% fewer agents who act according to the heuristic.¹⁶

The second example shows how an endogenous concern for misspecification helps firms to better choose the quantity to produce even if their models fail to capture how prices respond to increased production.

Example 2 (Belief in Price Taking, Sobel, 1984). A firm chooses a quantity to produce $a \in A$ at cost $C(a)$. The average price the firm charges is $P = \mathcal{P}(a) + \varepsilon_{1,a}$, where $\mathcal{P} : \mathbb{R} \rightarrow \mathbb{R}$ is a strictly decreasing (inverse) demand function. Here, $y = P$, and the payoff is $u(a, y) = aP - C(a)$. The firm believes they are price taker and that the average price they can charge does not depend on their production, i.e., $P = \theta + \varepsilon_{2,a}$.

Here, $Q(a) = \{q^{P(a)}\}$, i.e., the firm estimates that the price they could charge for the $a + 1$ -th quantity equals the (higher) average price charged for the first a units. Therefore, by Proposition 1, in any B -NE, the firm produces more than the optimum, and in every c -robust equilibrium, this

¹⁵To see this in action concretely, consider the following specific values, $A = \{1, 2\}$, $\tau(1) = 1/4$, $\tau(2) = 1/2$, $C(1) = 0$, $C(2) = 2/5$, $c = 10$, $\sigma_1^2 = 1$, $\sigma_2^2 = 2$. In this case, the expected utility of action 1 and 2 are respectively $3/4$ and $3/5$, so low effort is the optimal best action. However, high effort is a Berk-Nash equilibrium because while playing 2 the incorrect model of the agent leads them to estimate a performance of $1/2$ for the low effort action (and being correct on the average performance of 2). The relative entropy between the true and best-fitting distribution is 0.1, see equation (11). Therefore, $a = 2$ is not a c -robust equilibrium because the average robust control evaluations for low and high effort at the belief induced by 2 are respectively -0.5 and -1.4 .

¹⁶In this discussion I 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 scheduling 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.

bias is reduced. When the firm is not perfectly able to explain the equilibrium data, they maintain a positive level of concern for misspecification, and they reduce their exposure to the uncertainty in the price they will be able to charge by lowering the produced quantity.

3.3.2 Bias Increase under Correlation Neglect, Esponda (2008)

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 3. A buyer with valuation $v \in V$ and a seller submit a (bid) price $a \in A$, and an ask price $s \in S \subseteq \mathbb{R}_+$, 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 . Here, $y = (v, s)$, and the payoff for the buyer is

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

The buyer mistakenly believes that the ask price and valuation are independent: $Q = \Delta(V) \times \Delta(S)$. Easy computations show that for every $a^* \in A$,¹⁷

$$Q(a^*) = \{q \in Q : \forall a \in A, \forall s \in S, q_a(s) = p_a^*(s), q_a(v) = p_a^*(v)\}.$$

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 extreme consequences of this bias are reinforced in a c -robust equilibrium. If, for some primitives, a complete unraveling of the market where the buyer bids 0 is obtained in a Berk-Nash equilibrium, it is obtained in a c -robust equilibrium. The agent still observes some correlation between valuations and ask prices, and this results in a positive $\min_{q \in Q} R(p_{a^*}^* || q_{a^*}) > 0$ and makes the agent less confident in their model. Since offering 0 gives a certain payoff, it is less sensitive to the misspecification concern, and, therefore, this positive concern makes market participation less desirable.

¹⁷See Appendix B.3.1 for the computations supporting the claims of this example.

4 Cycles

Part 3 of Theorem 2 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, by Theorem 2, the agent behavior cannot stabilize. I 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

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

be the set of models with the lowest average relative entropy from the actual DGP, 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 $\nu \in \Delta(Q)$ with

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

A mixed robust equilibrium allows multiple actions to be played but requires that the probability assigned to each action determines the beliefs and the concern for misspecification. Intuitively, suppose actions for which the models in Q 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 2. *For every $c > 0$ there exists a mixed c -robust equilibrium.*

When the action process almost surely does not converge, obtaining a weaker form of behavior stabilization is still possible, i.e., the convergence of the empirical action frequency, which allows for persistent changes in actions. Let $\alpha_t(h_t) \in \Delta(A)$ be the empirical action frequency in history h_t , defined as

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

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

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

Theorem 3. *If*

$$\Lambda(h_t) = \frac{LLR(h_t, Q)}{ct} \quad \forall t \in \mathbb{N}, \forall h_t \in \mathcal{H}_t$$

and α^ is a Λ -limit frequency, then α^* is a mixed c -robust equilibrium.*

To interpret Theorem 3, consider the case where α^* is supported over two actions a, a' such that Q explains very well the consequences of a , —i.e., $\min_{q \in Q} R(p_a^* || q_a)$ is low— but it explains poorly the consequences of a' —i.e., $\min_{q \in Q} R(p_{a'}^* || q_{a'})$ is high. Suppose also that a is a best reply to a high misspecification concern, while a' best replies to a low misspecification concern. Then, the agent oscillates between periods with great concern for misspecification, when they play a , and phases in which the excellent data fit leads them to experiment with action a' .

4.1 Application: Monetary Policy Cycles

Here I consider a monetary policy model taken from Cogley and Sargent (2005) and Sargent (2008) and, in particular, its adaptation in Battigalli et al. (2022).¹⁸ A central bank tries 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 $a = 0$. The central bank assigns a positive probability to models parametrized by the vector $\theta \in \Theta \subseteq \mathbb{R}^5$, 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 $\varepsilon_U, \varepsilon_\pi$ are independent, zero-mean random shocks normalized to have the same support $[-1, 1]$. Here $\theta_0 > 0$ is the natural unemployment level, $\theta_{1\pi} < 0$ is the impact of the actual inflation on unemployment, and $\theta_{1a} > 0$ is the impact of the planned inflation on unemployment, a reduced form of the fact that the market participants (partially) incorporate the central bank's actions in their inflation expectations. In particular, if $\theta_{1\pi} = -\theta_{1a}$, this is a Lucas-Sargent model with no (structural) exploitable employment-inflation trade-off. If $\theta_{1\pi} < -\theta_{1a}$, this is a Samuelson-Solow model with a structural exploitable employment-inflation trade-off.

The central bank's model is misspecified in that it misses the fact that an aggressive monetary

¹⁸Spiegler (2020) also considers the effect of a misspecified model in the context of Phillips curve estimation, with the main difference being that the misspecification is on the side of the market.

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 $\theta_2^* > 0$, $\theta_3^* > 0$, f_0 is the identity function, while $f_1 : [-1, 1] \rightarrow [-1, 1]$ is a continuous, strictly increasing, and odd function with $f_1(1) = 1$ that is strictly concave on $[0, 1]$, i.e., that amplifies the inflation-specific shocks.

The central bank is endowed with standard quadratic preferences:

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

Assumption 3. i) $(\theta_{1\pi}^* + \theta_{1a}^* + \theta_0^*)^2 + 1 < (\theta_0^*)^2$. ii) $\text{essinf}_{p_1^*} u(1, y) < \text{essinf}_{p_0^*} u(0, y)$. iii) $\theta^* \in \Theta$ and for all $\theta \in \Theta$, $(\theta_{1a}, \theta_2, \theta_3) = (\theta_{1a}^*, \theta_2^*, \theta_3^*)$, $\theta_0 > 0, \theta_{1\pi} < 0, \theta_0 + \theta_{1\pi} + \theta_{1a} > 0$.

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

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 best replies to a low misspecification concern and θ^* but induces a severe concern. Therefore, the behavior cannot stabilize consistently with the cyclical behavior of monetary policies documented in Sargent (2008). There is also some natural comparative statics in the extremal robust equilibria, as a larger exploitable trade-off between inflation and unemployment induces more time spent using an aggressive monetary policy.

5 Discussion

5.1 Forward-looking Agents

One key generalization to our model would be to allow for forward-looking agents. As for many nonSEU decision criteria, the main complication is that the most immediate extension of the criterion to forward-looking agents would induce dynamic inconsistencies under some information structures (see Example 4 in the Supplemental Material for a simple explicit example). One approach would be to directly impose a recursive formulation for the preferences, as in Maccheroni et al. (2006b) and Klibanoff et al. (2009). Since the decision criterion belongs to the variational class, it is known from Maccheroni et al. (2006b) that a recursive formulation can be obtained, and Cerreia-Vioglio et al. (2024) indeed characterize its recursive formulation in the high patience limit. A complementary approach does not impose recursivity and allows for dynamic inconsistency. However, analyzing an uncommitted, forward-looking, and sophisticated agent playing an intra-personal equilibrium with their future selves would require combining the insights of this paper with the approach developed in Battigalli et al. (2019). See also Cerreia-Vioglio et al. (2024) for weaker notions of dynamic consistency compatible with an evolving concern for misspecification.

5.2 Model Dependent Concern

Throughout the paper, a single number, λ , summarizes the agent’s concern for misspecification. An alternative is to have a collection $(\lambda_q(h_t))_{q \in Q}$ of misspecification concerns, each of them adapted to the fit of the specific q to the observed data, i.e., with $\lambda_q(h_t) = \frac{LLR(h_t, \{q\})}{c}$. In the axiomatic foundation of Section 2 of the Supplemental Material, I explicitly note that having a homogeneous λ is a consequence of the axiom “Uniform Misspecification Concern”.

Still, it seems more realistic that confidence in a model acts at the level of its defining assumption, not at the level of a particular estimate. To be concrete, I interpret the concern for misspecification as a concern for the linearity hypothesis on the tax schedule or for omitting a potentially relevant variable (e.g., monetary policy in another country) from the economy description employed by the central bank. Not as a collection $(\lambda_x)_{x \in [0,1]}$ of concerns for “ $x\%$ being the exact tax rate”. Therefore, the agent computes the fit of the best explanation under linearity or irrelevance of the foreign monetary policy, using a unique $LLR(a^t, y^t, Q)$ rather than $(LLR(a^t, y^t, \{q\}))_{q \in Q}$. Moreover, having $\lambda_q(h_t) = \frac{LLR(h_t, \{q\})}{c}$ would induce a “double counting of evidence”. Indeed,

if q explains the data h_t better than q' , i.e., $\frac{LLR(h_t, \{q\})}{LLR(h_t, \{q'\})} < 1$ this will already impact the future evaluations by making $\frac{\mu(q|h_t)}{\mu(q'|h_t)} > \frac{\mu(q)}{\mu(q')}$. Having also $\frac{\lambda_q(h_t)}{\lambda_{q'}(h_t)} = \frac{LLR(h_t, \{q\})}{LLR(h_t, \{q'\})} < 1$ would duplicate the penalization already present in Bayesian updating. In this regard, our approach maintains a separation between the confidence $\frac{1}{\lambda}$ *in the overall specification* and the relative belief in two specific data-generating processes $\frac{\mu(q)}{\mu(q')}$ *within* the specification.

Nevertheless, the results of the paper hold even letting $\lambda_q(h_t) = \frac{LLR(h_t, \{q\})}{c}$. That is a consequence of the double-counting issue highlighted above. For example, consider Theorem 3. The proof of the theorem shows that if the empirical action distribution converges to $\alpha \in \Delta(A)$, beliefs concentrate around $Q(\alpha)$. But asymptotically, all the structured models in $Q(\alpha)$ fit equally well p^* : $\lim_{t \rightarrow \infty} LLR(h_t, \{q\}) = \lim_{t \rightarrow \infty} LLR(h_t, Q)$ for all $q \in Q(\alpha)$. Evidence forces the models that are still relevant in the long run to share the same λ_q even if, in principle, they were allowed to be structured-model specific. Of course, this does not mean the two criteria are equivalent in the short run.

A similar consideration applies to a reasonable alternative rule for adjusting the misspecification concern, the posterior Bayes factor (Aitkin, 1991). This would require to put a prior $\nu \in \Delta(N(Q))$ also on the alternative set of models and have the concern for misspecification $\lambda^B(h_t) = -\log\left(\frac{\int_Q \prod_{\tau=1}^t \tilde{q}_{a_\tau}(y_\tau) d\mu(q|h_t)}{\int_{N(Q)} \prod_{\tau=1}^t \tilde{p}_{a_\tau}(y_\tau) d\nu(q|h_t)}\right)$. However, along every sequence of histories $(h_t)_{t \in \mathbb{N}}$, $\lim_{t \rightarrow \infty} (\lambda^B(h_t) - \lambda(h_t)) = 0$, and with this all the results of the paper extend verbatim.

A Appendix

By Assumption 1, there exists $K \in \mathbb{R}_{++}$ such that for all $a \in A$

$$-\ln \tilde{q}_a \leq K \quad \forall q \in Q$$

holds p_a^* -a.s. Throughout the Appendix, the symbol K will denote such a strictly positive real number. For every $t \in \mathbb{N}$ and history $h_t = (a^t, y^t) \in \mathcal{H}$ let $p^{h_t} \in \Delta(Y)^A$ be the action contingent (finite support) probability measure over outcomes corresponding to the empirical frequency: for all $a \in A$ such that $\sum_{\tau=1}^t \mathbb{I}_{\{a\}}(a_\tau) > 0$, $p_a^{h_t}(C) = \frac{\sum_{\tau=1}^t \mathbb{I}_{\{(a, y): y \in C\}}(a_\tau, y_\tau)}{\sum_{\tau=1}^t \mathbb{I}_{\{a\}}(a_\tau)}$ for all $C \in \mathcal{B}(Y)$ and $p_a^{h_t} = \delta_{\bar{y}}$ for some arbitrary fixed $\bar{y} \in Y$ if $\sum_{\tau=1}^t \mathbb{I}_{\{a\}}(a_\tau) = 0$. For all $b \in A$ let Π^b the policy that prescribes

b at every period. Let

$$M_\varepsilon(\alpha) = \left\{ \nu \in \Delta(Q) : \int_Q \sum_{a \in A} \alpha(a) R(p_a^* || q_a) d\nu(q) \leq \varepsilon + \min_{q \in Q} \sum_{a \in A} \alpha(a) R(p_a^* || q_a) \right\}$$

and $Q^\varepsilon(\alpha) := \{q \in Q : \exists q' \in Q(\alpha) \cap B_\varepsilon(q)\}$. Together, the next three lemmas show that an average robust control evaluation converges to a subjective expected utility evaluation (resp. maxmin evaluation) as λ tends to 0 (resp. to $+\infty$), generalizing previous results in the decision-theoretic literature in three dimensions (i) The function evaluated is not a finite range one (ii) There is a possibly infinite set of robust control models over which the average is taken (iii) The weights in the average are also allowed to change as λ changes.

Lemma 1. *For every $a \in A$,*

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

Proof. By Assumption 1 (i), $\text{supp} q_a = \text{supp} p_a^*$ for every $q \in Q$. Let $\hat{y} \in \text{argmin}_{y \in \text{supp} p_a^*} u(a, y)$. If $\max_{y \in Y} u(a, y) = u(a, \hat{y})$ the statement is trivially true, so suppose $\max_{y \in Y} u(a, y) > u(a, \hat{y})$. Fix $\bar{\varepsilon} \in \left(0, \frac{\max_{y \in Y} u(a, y) - u(a, \hat{y})}{2}\right)$. Since $u(a, \cdot)$ is continuous, there exists ε such that $y \in B_\varepsilon(\hat{y})$ implies $u(a, y) \leq u(a, \hat{y}) + \bar{\varepsilon}$. Then, by Proposition 1.4.2 in Dupuis and Ellis (2011), for all $q \in Q$

$$u(a, \hat{y}) \leq \min_{p_a \in \Delta(Y)} \int_Y u(a, y) dp_a + \frac{R(p_a || q_a)}{k} \leq -\log \left(\frac{\exp(-k(u(a, \hat{y}) + \bar{\varepsilon})) \inf_{\hat{q} \in Q} \hat{q}_a(B_\varepsilon(\hat{y}))}{+ (1 - \inf_{\hat{q} \in Q} \hat{q}_a(B_\varepsilon(\hat{y}))) e^{-k \max_{y \in Y} u(a, y)}} \right) / k.$$

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 q this proves the desired uniformity of the convergence. \blacksquare

Lemma 2. *For every $a \in A$, $q \in Q$, and $(q_n, \lambda_n)_{n \in \mathbb{N}} \in (Q \times \mathbb{R}_{++})^{\mathbb{N}}$ with $\lim_{n \rightarrow \infty} (q_n, \lambda_n) = (q, 0)$, $\lim_{n \rightarrow \infty} \min_{p_a \in \Delta(Y)} \left(\mathbb{E}_{p_a} [u(a, y)] + \frac{R(p_a || q_{a, n})}{\lambda_n} \right) = \mathbb{E}_{q_a} [u(a, y)]$.*

Lemma 3. *1. For every $a \in A$, the function $G : \Delta(Q) \times \mathbb{R}_+ \rightarrow \mathbb{R}$ defined by*

$$G(\nu, \lambda) := \int_Q \min_{p_a \in \Delta(Y)} \left(\mathbb{E}_{p_a} [u(a, y)] + \frac{R(p_a || q_a)}{\lambda} \right) d\nu(q) \quad \forall \nu \in \Delta(Q), \forall \lambda \in \mathbb{R}_{++}$$

and $G(\nu, 0) := \int_Q \mathbb{E}_{q_a} [u(a, y)] d\nu(q)$ for all $\nu \in \Delta(Q)$ is continuous. Moreover, the set of functions $\{G(\nu, \cdot)\}_{\nu \in \Delta(Q)}$ is equicontinuous at 0.

2. The correspondence $BR^{(\cdot)}(\cdot) : \mathbb{R}_+ \times \Delta(Q) \rightrightarrows A$ where $BR^0(\nu) := BR^{Seu}(\nu)$ for all $\nu \in \Delta(Q)$ is upper hemicontinuous.

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

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

Since Y is compact and u is continuous, for all $\lambda \in \mathbb{R}_{++}$ and $q \in Q$, the RHS belongs to $[\min_{y \in Y} u(a, y), \max_{y \in Y} u(a, y)]$. For every $\lambda \in \mathbb{R}_{++}$, $\exp(-\lambda u(a, \cdot))$ is a continuous and bounded function that is bounded away from 0. Therefore, $q \mapsto \int_Y \exp(-\lambda u(a, y)) dq_a(y)$ is continuous by definition of the weak convergence of measures, and F is continuous by Lemma 2 (at $\lambda = 0$) and Theorem 15.7.3 in Kallenberg (1973) (at $\lambda \neq 0$).

Let $(\nu_n, \lambda_n)_{n \in \mathbb{N}} \in \Delta(Q) \times \mathbb{R}_{++}$ be a convergent sequence with limit (ν, λ) . Suppose first that $\lambda > 0$. Then $\lim_{n \rightarrow \infty} \int_Q \frac{\log(\int_Y \exp(-\lambda_n u(a, y)) dq_a(y))}{-\lambda_n} d\nu_n(q) = \int_Q \frac{\log(\int_Y \exp(-\lambda u(a, y)) dq_a(y))}{-\lambda} d\nu(q)$ by Theorem 15.7.3 in Kallenberg (1973) and the joint continuity of F established above. Next, suppose that $\lambda = 0$. Then $\lim_{n \rightarrow \infty} \int_Q \frac{\log(\mathbb{E}_{q_a}[\exp(-\lambda_n u(a, y))])}{-\lambda_n} d\nu_n(q) = \int_Q \mathbb{E}_{q_a} [u(a, y)] d\nu(q)$ again by Theorem 15.7.3 in Kallenberg (1973) and the joint continuity of F established above.

Next, suppose that $\{G(\nu, \cdot)\}_{\nu \in \Delta(Q)}$ is not equicontinuous at 0. Then, there exists ε and $(\nu_n)_{n \in \mathbb{N}}$ such that $|G(\nu_n, 1/n) - G(\nu_n, 0)| > \varepsilon$ for all $n \in \mathbb{N}$. Since Q is compact, ν_n can be taken to be convergent to some ν' . But from the first part of the statement, it would follow that $\lim_{n \rightarrow \infty} G(\nu_n, 1/n) = G(\nu', 0) = \lim_{n \rightarrow \infty} G(\nu_n, 0)$, a contradiction.

(2) Follows by (1) and the Maximum Theorem. ■

Lemma 4. For all $\Pi \in A^{\mathcal{H}}$, $\mathbb{P}_{\Pi} \left(\left\{ (a_i, y_i)_{i \in \mathbb{N}} \in (A \times Y)^{\mathbb{N}} : \forall (t, q) \in \mathbb{N} \times Q, \ln \frac{1}{\tilde{q}_{a_t}(y_t)} \leq K \right\} \right) = 1$.

Proof of Lemma 4. By Assumption 1 for all $a \in A$, for every $t \in \mathbb{N}$,

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

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

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

proving the statement. ■

Lemma 5. For every $a \in A$, if $(q_n, p_a^n)_{n \in \mathbb{N}} \in (Q \times \Delta(Y))^{\mathbb{N}}$ is such that $\lim_{n \rightarrow \infty} (q_n, p_a^n)_{n \in \mathbb{N}} = (q', \bar{p}_a)$ and $\text{supp} p_a^n \subseteq \{y \in Y : \max_{q \in Q} -\ln \tilde{q}_a(y) \leq K\}$ then $\lim_{n \rightarrow \infty} -\int_Y \log(\tilde{q}_{a,n}(y)) dp_a^n(y) = -\int_Y \log(\tilde{q}'_a(y)) d\bar{p}_a(y)$.

Lemma 6. Let $\alpha^* \in \Delta(A)$, $(t_n)_{n \in \mathbb{N}}$ be a subsequence of \mathbb{N} , and $(a_t, y_t)_{t \in \mathbb{N}} \in (A \times Y)^{\mathbb{N}}$ be such that $\sup_{q \in Q, t \in \mathbb{N}} -\ln \tilde{q}_{a_t}(y_t) \leq K$. For every $n \in \mathbb{N}$, set $h_{t_n} = (a^{t_n}, y^{t_n})$, and let $q(h_{t_n})$ and $r(h_{t_n})$ be two arbitrary elements of $\text{argmax}_{q \in Q} \prod_{\tau=1}^{t_n} \tilde{q}_{a_\tau}(y_\tau)$ and $\text{argmax}_{p \in N(Q)} \prod_{\tau=1}^{t_n} \tilde{p}_{a_\tau}(y_\tau)$, respectively. If

$$\lim_{n \rightarrow \infty} \left(\alpha_{t_n}(h_{t_n}), (p_a^{h_{t_n}})_{a \in \text{supp} \alpha^*} \right) = \left(\alpha^*, (p_a^*)_{a \in \text{supp} \alpha^*} \right)$$

then

$$\lim_{n \rightarrow \infty} \frac{LLR(h_{t_n}, Q)}{t_n} = \lim_{n \rightarrow \infty} \frac{\sum_{a \in A} \sum_{\tau=1}^{t_n} \mathbb{I}_{\{a\}}(a_\tau) \int_Y \log \left(\frac{\tilde{r}_a(h_{t_n})(y)}{\tilde{q}_a(h_{t_n})(y)} \right) dp_a^{h_{t_n}}(y)}{t_n} = \min_{q \in Q} \sum_{a \in A} \alpha^*(a) R(p_a^* || q_a).$$

Proof. By assumption of the lemma, for all $n \in \mathbb{N}$,

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

By Assumption 2 (i),

$$\sum_{a \in A} \sum_{\tau=1}^{t_n} \mathbb{I}_{\{a\}}(a_\tau) \int_Y -\log(\tilde{r}_a(h_{t_n})(y)) dp_a^{h_{t_n}}(y) / t_n \leq 0 \quad \forall n \in \mathbb{N}.$$

By Assumption 2 (ii) and the Arzela-Ascoli Theorem, there exists $K' \in \mathbb{R}_{++}$ such that for all $a \in A$,

$$-\log(\tilde{r}_a(h_{t_n})(y)) \geq -K', \tag{6}$$

p_a^* -almost surely. Therefore, for all $n \in \mathbb{N}$, $\sum_{a \in A} \sum_{\tau=1}^{t_n} \mathbb{I}_{\{a\}}(a_\tau) \int_Y -\log(\tilde{r}_a(h_{t_n})(y)) dp_a^{h_{t_n}}(y) / t_n \subseteq [-K', 0]$. I now show that every convergent subsequence converges to 0. Indeed, take any such subsequence of periods $(t_{n_i})_{i \in \mathbb{N}}$ and, possibly restricting to a further subsequence, suppose $\tilde{r}_a(h_{t_{n_i}})$

converges to some \tilde{r}_a for every $a \in \text{supp}\alpha^*$ (this can be done by Assumption 2 (ii)). Then

$$\begin{aligned} 0 &\leq \sum_{a \in A} \alpha^*(a) \int_Y -\log(\tilde{r}_a(y)) dp_a^*(y) \leq \sum_{a \in A} \alpha^*(a) \liminf_{n \rightarrow \infty} \int_Y -\log(\tilde{r}_a(h_{t_{n_i}})(y)) dp_a^{h_{t_{n_i}}}(y) \\ &= \liminf_{i \rightarrow \infty} \sum_{a \in A} \sum_{\tau=1}^{t_{n_i}} \mathbb{I}_{\{a\}}(a_\tau) \int_Y -\log(\tilde{r}_a(h_{t_{n_i}})(y)) dp_a^{h_{t_{n_i}}}(y) / t_{n_i} \end{aligned}$$

where the first inequality follows from Gibbs inequality and the second since by equation (6) I can apply Lemma 3.2 in Serfozo (1982). Therefore,

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

Since this holds for every converging subsequence, it holds for the original sequence. So

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

Therefore the result follows from Lemma 5 and Theorem 17.31 in Aliprantis and Border (2013). \blacksquare

Lemma 7. For every $t \in \mathbb{N}$, $\Pi \in A^{\mathcal{H}}$ and for \mathbb{P}_Π almost every $h_t = (a^t, y^t) \in \mathcal{H}_t$, if $q' \in \text{argmax}_{q \in Q} \prod_{\tau=1}^t \tilde{q}_{a_\tau}(y_\tau)$, and $p \in \text{argmax}_{r \in N(Q)} \prod_{\tau=1}^t \tilde{r}_{a_\tau}(y_\tau)$ then

$$LLR(h_t, Q) = \sum_{a \in A} \sum_{\tau=1}^t \mathbb{I}_{\{a\}}(a_\tau) \int_Y \log\left(\frac{dp_a(y)}{dq'_a(y)}\right) dp_a^{h_t}(y).$$

Proof of Theorem 1. I start observing that by Lemma 4, $\max_{q \in Q, t \in \mathbb{N}} -\ln \tilde{q}_{\mathbf{a}_t}(\mathbf{y}_t) \leq K$, $\mathbb{P}_{\Pi^{\alpha^*}}$ -a.s. This will allow us to invoke Lemma 6 repeatedly.

Proof of Part 1. Let (u, A, Y) be a decision problem and $\varepsilon \in \mathbb{R}_{++}$. Let $\hat{P} \subseteq \Delta(Y)^A$ be the set of p^* that satisfy Assumption 1 jointly with Q . I start by showing that there exists $c \in \mathbb{R}_{++}$

such that the rule of equation (2) is ε -safe and ε -self-confirming under correct specification. This is done by first deriving a $c \in \mathbb{R}_{++}$ such that ε -safety is satisfied, and then showing that there exists a δ that delivers ε -self-confirming under correct specification.

Safety is trivially satisfied by every policy if $\max_{a \in A} \min_{y \in Y} u(a, y) = \min_{a \in A, y \in Y} u(a, y)$, so in that case pick an arbitrary $c \in \mathbb{R}_{++}$. Suppose $\max_{a \in A} \min_{y \in Y} u(a, y) > \min_{a \in A, y \in Y} u(a, y)$. Define $\underline{A}(p^*) := \left\{ a' \in A : \max_{a \in A} \min_{y \in Y} u(a, y) > \mathbb{E}_{p_{a'}} [u(a', y)] + \frac{\varepsilon}{2} \right\}$ for all $p^* \in \hat{P}$.

Claim 1. *There exists $\varphi^* > 0$ such that for every $\Pi \in A^{\mathcal{H}}$ and $p^* \in \hat{P}$,*

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

That is, almost surely the payoff is at most ε -lower than the safe guarantee if the actions whose objective expected performance is lower than the guarantee are played sufficiently rarely (i.e., each of them has a frequency lower than φ^*).

Proof. Consider the stochastic process defined by $\mathbf{X}_t = u(\Pi(\mathbf{h}_{t-1}), \mathbf{y}_t) - \mathbb{E}_{p_{\Pi(\mathbf{h}_{t-1})}^*} [u(\Pi(\mathbf{h}_{t-1}), y)]$, for all $t \in \mathbb{N}$ with the sequence of sigma-algebras $(\mathcal{F}_t)_{t \in \mathbb{N}}$ generated by the stochastic process of histories $(\mathbf{h}_t)_{t \in \mathbb{N}}$. The stochastic process is a martingale difference sequence, as u is continuous in y on the compact Y , so $\mathbb{E}[|\mathbf{X}_t|] < \infty$ and $\mathbb{E}[\mathbf{X}_t | \mathcal{F}_{t-1}] = 0$. By the SLLN for martingale difference sequences, $\lim_{n \rightarrow \infty} \frac{\sum_{t=1}^n \mathbf{X}_t}{n} = 0$, \mathbb{P}_{Π} -a.s. so that

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

and the claim follows from setting $\frac{\varepsilon}{2(\max_{\bar{a} \in A} \min_{y \in Y} u(\bar{a}, y) - \min_{a \in A, y \in Y} u(a, y))|A|} = \varphi^*$. \square

Claim 2. *There exists $\bar{\lambda} \in \mathbb{R}_{++}$ such that if $\lambda \geq \bar{\lambda}$ then for every $p^* \in \hat{P}$, $a' \in \underline{A}(p^*)$, $\nu \in \Delta(Q)$, $a' \notin BR^\lambda(\nu)$.*

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

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

Claim 3. *There exists $J \in (0, 1)$ such that for every $p^* \in \hat{P}$, $a' \in \underline{A}(p^*)$, $\mu \in \Delta(Q)$, and $\lambda \in \mathbb{R}_+$,*

$$\mu(\{q \in Q : R(p_{a'}^* || q_{a'}) > J\}) \leq J \implies a' \notin BR^\lambda(\mu). \quad (7)$$

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

Proof. Observe that given Claim 2, the statement immediately holds for $\lambda > \bar{\lambda}$ for any $J \in (0, 1)$. Suppose by contradiction that equation (7) does not hold. This means that there exist a convergent $(p_n^*, \mu_n, \lambda_n)_{n \in \mathbb{N}} \in \hat{P} \times \Delta(Q) \times [0, \bar{\lambda}]$ and $a' \in A$ with $\mu_n(\{q \in Q : R(p_{a',n}^* || q_{a'}) > \frac{1}{n}\}) \leq \frac{1}{n}$, $a' \in \underline{A}(p_n^*)$, and $a' \in BR^{\lambda_n}(\mu_n)$. By the lower semicontinuity of R and the fact that $R(p_{a'} || q_{a'}) = 0$ if and only if $p_{a'} = q_{a'}$, (see, e.g., Lemma 1.4.3 in Dupuis and Ellis (2011)), as well as Lemma 3 this in turn implies that there exists $q \in Q$ with $a' \in \underline{A}(q)$ and $\mathbb{E}_{q_{a'}}[u(a', y)] \geq \max_{\bar{a} \in A} \min_{y \in Y} u(\bar{a}, y)$, a contradiction. \square

Let $c = \frac{J\varphi^*}{4\bar{\lambda}}$, and take an arbitrary $p^* \in \hat{P}$. If for all $a' \in \underline{A}(p^*)$, $\limsup_{t \rightarrow \infty} \alpha_t(\mathbf{h}_t)(a') < \varphi^*$, \mathbb{P}_Π -a.s., ε -safety follows by Claim 1. Suppose by contradiction that there is an action $a' \in \underline{A}(p^*)$ with $\limsup_{t \rightarrow \infty} \alpha_t(\mathbf{h}_t)(a') \geq \varphi^*$ on a set E with \mathbb{P}_Π strictly positive probability. Therefore, on all the infinite histories in E there must be a subsequence of periods $(t_n)_{n \in \mathbb{N}}$ in which both $\alpha_{t_n}(h_{t_n})(a') \geq \varphi^*/2$ and $a_{t_n+1} = a'$ for all $n \in \mathbb{N}$. By the compactness of $\Delta(A)$, I can take this subsequence to be also such that $\lim_{n \rightarrow \infty} \alpha_{t_n} = \bar{\alpha}$. By Claim 3, it must be the case that

$$\mu(\{q \in Q : R(p_{a'}^* || q_{a'}) > J\} | h_{t_n}) \geq J \quad \forall n \in \mathbb{N}.$$

But then,

$$\lim_{n \rightarrow \infty} (\Lambda(h_{t_n}))_{n \in \mathbb{N}} = \lim_{n \rightarrow \infty} \frac{LLR(h_{t_n}, Q)}{ct_n} = \min_{q \in Q} \sum_{a \in A} \bar{\alpha}(a) R(p_a^* || q_a) / c \geq \frac{\varphi^* J}{4c} \geq 2\bar{\lambda} \quad \mathbb{P}_\Pi\text{-a.s.}$$

where the second equality follows by Lemma 6, whose assumptions are in turn satisfied by Theorem

11.4.1 in Dudley (2018).¹⁹ But, by Claim 2, for sufficiently large n this leads to the contradiction with $a_{t_n+1} = a'$.

I now move to prove ε -self-confirming under correct specification. I start with a few considerations for a fixed DGP $p^* \in \hat{P}$. The next claim parallels Claim 1.

Claim 4. Let $\Pi \in A^{\mathcal{H}}$ and $p^* \in \hat{P}$,

$$\mathbb{P}_{\Pi} \left(\begin{array}{c} \left\{ \liminf_{t \rightarrow \infty} \frac{\sum_{i=1}^t u(\mathbf{a}_i, \mathbf{y}_i)}{t} \leq G_{\varepsilon}(p^*) - \varepsilon \right\} \\ \cap \left\{ \limsup_{t \rightarrow \infty} \alpha_t(\mathbf{h}_t)(a') < \frac{\varepsilon}{2|A|(\max_{a \in A, y \in Y} u(a, y) - \min_{a \in A, y \in Y} u(a, y))}, \forall a' \notin SCE_{\varepsilon}(p^*) \right\} \end{array} \right) = 0.$$

Proof. Consider the stochastic process defined by $\mathbf{X}_t = u(\Pi(\mathbf{h}_{t-1}), \mathbf{y}_t) - \mathbb{E}_{p_{\Pi(\mathbf{h}_{t-1})}^*} [u(\Pi(\mathbf{h}_{t-1}), y)]$, for all $t \in \mathbb{N}$ with the sequence of sigma-algebras $(\mathcal{F}_t)_{t \in \mathbb{N}}$ generated by the stochastic process of histories $(\mathbf{h}_t)_{t \in \mathbb{N}}$. The stochastic process is a martingale difference sequence, as u is continuous in y on the compact Y , so $\mathbb{E}[|\mathbf{X}_t|] < \infty$ and $\mathbb{E}[\mathbf{X}_t | \mathcal{F}_{t-1}] = 0$. By the SLLN for martingale difference sequences, $\lim_{n \rightarrow \infty} \frac{\sum_{i=1}^n \mathbf{X}_i}{n} = 0$, \mathbb{P}_{Π} -a.s. so that

$$\begin{aligned} \liminf_{t \rightarrow \infty} \frac{\sum_{i=1}^t u(\mathbf{a}_i, \mathbf{y}_i)}{t} &= \liminf_{t \rightarrow \infty} \sum_{i=1}^t \frac{\mathbf{X}_i + \mathbb{E}_{p_{\mathbf{a}_i}^*} [u(\mathbf{a}_i, \cdot)]}{t} \\ &\geq \left(1 - \sum_{a \notin SCE_{\varepsilon}(p^*)} \limsup_{t \rightarrow \infty} \alpha_t(\mathbf{h}_t)(a) \right) G_{\varepsilon}(p^*) + \sum_{a \notin SCE_{\varepsilon}(p^*)} \limsup_{t \rightarrow \infty} \alpha_t(\mathbf{h}_t)(a) \min_{a \in A, y \in Y} u(a, y) \\ &\geq \left(1 - |A| \max_{a \notin SCE_{\varepsilon}(p^*)} \limsup_{t \rightarrow \infty} \alpha_t(\mathbf{h}_t)(a) \right) G_{\varepsilon}(p^*) + |A| \max_{a \notin SCE_{\varepsilon}(p^*)} \limsup_{t \rightarrow \infty} \alpha_t(\mathbf{h}_t)(a) \min_{a \in A, y \in Y} u(a, y) \\ &\geq \left(1 - |A| \max_{a \notin SCE_{\varepsilon}(p^*)} \limsup_{t \rightarrow \infty} \alpha_t(\mathbf{h}_t)(a) \right) G_{\varepsilon}(p^*) \\ &\quad + |A| \max_{a \notin SCE_{\varepsilon}(p^*)} \limsup_{t \rightarrow \infty} \alpha_t(\mathbf{h}_t)(a) [G_{\varepsilon}(p^*) - \max_{a \in A, y \in Y} u(a, y) + \min_{a \in A, y \in Y} u(a, y)]. \end{aligned}$$

□

Also observe that $\limsup_{t \rightarrow \infty} \alpha_t(\mathbf{h}_t)(\underline{a}) > \varphi^{**} := \frac{\varepsilon}{2|A|(\max_{a \in A, y \in Y} u(a, y) - \min_{a \in A, y \in Y} u(a, y))}$ means that infinitely many times

$$\alpha_t(\mathbf{h}_t)(\underline{a}) > \varphi^{**}/2 \text{ and } \underline{a} \in BR^{\Lambda(\mathbf{h}_t)}(\mu(\cdot | \mathbf{h}_t)). \quad (8)$$

¹⁹See Remark 3 in the Supplemental Material for a subtlety in the use of Theorem 11.4.1 in Dudley (2018).

I now show that there exists a $\delta > 0$ such that this cannot happen with positive probability under any $p^* \in \hat{P}$ for any $a \notin SCE_\varepsilon(p^*)$ if $\min_{q \in Q} \max_{a \in A} R(p_a^* || q_a) < \delta$.

By Lemma 3, there exists $\delta \in (0, \varepsilon \varphi^{**}/4)$ such that for all $p^* \in \hat{P}$, $a', a'' \in A$ and $\alpha \in \Delta(A)$ with $\alpha(a') \geq \varphi^{**}/2$,

$$\min_{q \in Q} \max_{a \in A} R(p_a^* || q_a) < \delta \quad (9)$$

implies that for $\nu \in \Delta(M_{\varepsilon \varphi^{**}/4}(\alpha))$ and $\lambda \in [0, 2\delta/c]$ (where recall that c has been set to $\frac{J\varphi^*}{4\lambda}$),²⁰

$$\begin{aligned} & \int_Q \min_{p \in \Delta(Y)} \left(\mathbb{E}_p[u(a', y)] + \frac{R(p_{a'} || q_{a'})}{\lambda} \right) d\nu(q) - \int_Q \min_{p \in \Delta(Y)} \left(\mathbb{E}_p[u(a'', y)] + \frac{R(p_{a''} || q_{a''})}{\lambda} \right) d\nu(q) \\ & \leq \max_{\nu' \in \Delta(\{q \in Q: R(p_{a'} || q_{a'}) \leq \varepsilon\})} \left\{ \int_Q \mathbb{E}_{q_{a'}}[u(a', y)] d\nu'(q) - \int_Q \mathbb{E}_{q_{a''}}[u(a'', y)] d\nu'(q) \right\} + \frac{\varepsilon}{4}. \end{aligned} \quad (10)$$

In words, given that action a' is played with sufficiently high probability (i.e., $\alpha(a') \geq \varphi^{**}/2$), the choice of a sufficiently small δ can make one of the models sufficiently close to the true DGP (i.e., equation (9)) and λ sufficiently small (i.e., $\lambda \leq 2\delta$) such that the (average robust control) evaluation with respect to the $\varepsilon \varphi^{**}/4$ best models (i.e., $\nu \in \Delta(M_{\varepsilon \varphi^{**}/4}(\alpha))$) of action a' with respect to a'' cannot be more than $\frac{\varepsilon}{4}$ more optimistic than the (SEU) evaluation under the models that come ε close to explain the consequence of action a' .

Observe that if $a' \notin SCE_\varepsilon(p^*)$, independently of the p^* , the RHS of equation (10) is smaller than $-\frac{3\varepsilon}{4}$ for some $a'' \in A$. But then, by Claim 4, Λ is ε -self-confirming with this δ equation (8), since by Theorem 1 in Esponda et al. (2021a) beliefs concentrate \mathbb{P}_Π almost surely on $M_{\varepsilon \varphi^{**}/4}(\alpha(h_t))$, and by Lemma 6, $\lambda \leq 2\delta/c$ at histories where $\alpha(a')(h_t) \geq \varphi^{**}/2$.²¹

Proof of Part 2. I show that there is a decision problem (u, A, Y) such that if the concern for misspecification of the agent is such that $\Lambda(\mathbf{h}_t) = o\left(\frac{LLR(\mathbf{h}_t, Q)}{t}\right)$, \mathbb{P}_Π -a.s. or $o(\Lambda(\mathbf{h}_t)) = \frac{LLR(\mathbf{h}_t, Q)}{t}$, \mathbb{P}_Π -a.s. then the decision rule cannot be both $\frac{1}{20}$ -safe and $\frac{1}{20}$ -self-confirming. Suppose that $A = \{1, -1, 0\}$ and $Y = \{-1, 1\}$. The utility function is $u(a, y) = ay$. Each model q considered by the agent is described by $q_a(1)$ for some arbitrary $a \in A$. Let $Q = \{0.9, 0.4\}$, $\mu(0.9) = \frac{1}{2} = \mu(0.4)$,

²⁰To see this, observe that

$$R(p_{a'}^* || q_{a'}) > \varepsilon \implies \sum_{a \in A} \alpha(a) R(p_a^* || q_a) > \varepsilon \varphi^{**}/2 \implies q \notin M_{\varepsilon \varphi^{**}/2-\delta}(\alpha) \implies q \notin M_{\varepsilon \varphi^{**}/4}(\alpha).$$

²¹Once again, the assumptions of Lemma 6 are satisfied by Theorem 11.4.1 in Dudley (2018). See also the discussion in Remark 3.

and Let $N(Q) = [0, 1]$, i.e., the unstructured models include all the action-independent DGPs.

I first show that if $\Lambda(\mathbf{h}_t) = o\left(\frac{LLR(\mathbf{h}_t, Q)}{t}\right)$, \mathbb{P}_Π -a.s. then the decision rule is not $\frac{1}{20}$ -safe. Let $p_a^*(1) = 0.6$. So, $\max_{\bar{a} \in A} \min_{y \in Y} u(\bar{a}, y) = \min_{y \in Y} u(0, y) = 0$. However, by the Strong Law of Large Numbers it follows that \mathbb{P}_Π -a.s. $\lim_{t \rightarrow \infty} (\sum_{\tau=1}^t \mathbb{I}_{\{1\}}(\mathbf{y}_\tau)) / t = 0.6$. Therefore, by Lemma 6,

$$\lim_{t \rightarrow \infty} \frac{LLR(\mathbf{h}_t, Q)}{t} = R(0.6||0.4) \quad \mathbb{P}_\Pi\text{-a.s.}$$

and so $\lim_{t \rightarrow \infty} \Lambda(\mathbf{h}_t) = 0$, \mathbb{P}_Π -a.s. Moreover, for the constant function $\phi(\varepsilon) = \frac{1}{2}$ for all $\varepsilon \in \mathbb{R}_{++}$ the prior is ϕ -positive on Q in the sense of Fudenberg et al. (2023), and by their Lemma 1, $\mu(0.4|\mathbf{h}_t) \rightarrow 1$, \mathbb{P}_Π -a.s. But then by the upper hemicontinuity of $BR^{(\cdot)}(\cdot)$ established in Lemma 3

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

proving the desired result.

Finally, I show that if the concern for misspecification of the agent is such that $o(\Lambda(\mathbf{h}_t)) = \frac{LLR(\mathbf{h}_t, Q)}{t}$, \mathbb{P}_Π -a.s. then the decision rule is not $\frac{1}{20}$ -self-confirming. Let Π be a Λ -optimal policy, $\delta \in (0, \frac{1}{20})$, and $p_a^*(1) = 0.9 + \delta$ for all $a \in A$, so that $SCE_{1/20}(p^*) = \{1\}$, and $G_{1/20}(p^*) = 0.2$.

Let $\bar{\lambda}$ be such that $\{0\} = BR^\lambda(\mu)$ for all $\lambda \geq \bar{\lambda}$ and $\mu \in \Delta(Q)$. Such a $\bar{\lambda}$ exists by Lemma 1. By Lemma 6 $\lim_{t \rightarrow \infty} LLR(h_t, Q) / t \geq R(0.9 + \delta || 0.9) / 2$, \mathbb{P}_π almost surely so that $\Lambda(h_t)$ is diverging to $+\infty$ and the result follows. \blacksquare

Proof of Theorem 2. I start observing that by Lemma 4, $\max_{q \in Q, t \in \mathbb{N}} -\ln \tilde{q}_{\mathbf{a}_t}(\mathbf{y}_t) \leq K$, $\mathbb{P}_{\Pi^{a^*}}$ -a.s. This will allow us to invoke Lemma 6 in all the various cases.

1) Suppose by contradiction that a^* is a Λ -limit action but is not a B-NE. Thus, since for every policy $\Pi \in A^\mathcal{H}$

$$\mathbb{P}_\Pi[\text{sup}\{\mathbf{t}: \mathbf{a}_\mathbf{t} \neq a^*\} < \infty] \leq \sum_{t=0}^{\infty} \sum_{h_t \in \mathcal{H}_t} \mathbb{P}_\Pi[a^* = \Pi(h_\tau), \forall \tau \geq t | h_t] \mathbb{P}_\Pi[h_t],$$

there are a Λ -optimal policy $\tilde{\Pi} \in A^\mathcal{H}$, $t \in \mathbb{N}_0$, and $h_t \in \mathcal{H}_t$ with $\mathbb{P}_{\tilde{\Pi}}[h_t] > 0$ such that with positive probability $\tilde{\Pi}$ prescribes a^* after h_t in every future period. Define $\nu = \mu(\cdot | h_t)$, and notice that by Assumption 1 (i) $\text{supp } \nu = \text{supp } \mu = Q$. As the evolution of beliefs and misspecification concern under Π^{a^*} , i.e., the policy that plays a^* in every period, is the same as under $\tilde{\Pi}$ for every history

where the agent continues to play a^* ,

$$\mathbb{P}_{\tilde{\Pi}}[a^* = \tilde{\Pi}(\mathbf{h}_\tau), \forall \tau > t | h_t] > 0 \implies \mathbb{P}_{\Pi^{a^*}}[a^* \in BR^{\Lambda(h_t, \mathbf{h}_\tau)}(\nu(\cdot | \mathbf{h}_\tau)), \forall \tau > t] > 0.$$

I now show that the latter equals zero, which establishes that a^* cannot be a Λ -limit action.

Since Y is a compact metric space, it is separable, and thus, by Varadarajan Theorem $\lim_{\tau \rightarrow \infty} p_{a^*}^{\mathbf{h}_\tau} = p_{a^*}^*$, $\mathbb{P}_{\Pi^{a^*}}$ -a.s. Then, by Lemma 6 and equation (4), $\lim_{\tau \rightarrow \infty} \Lambda(h_t, \mathbf{h}_\tau) = 0$, $\mathbb{P}_{\Pi^{a^*}}$ -a.s. By Assumption 1 (ii), the assumptions of Berk (1966), page 54, are satisfied, and for every $\varepsilon \in \mathbb{R}_{++}$, $\nu(Q^\varepsilon(a^*) | \mathbf{h}_\tau) \rightarrow 1$, $\mathbb{P}_{\Pi^{a^*}}$ -a.s. Therefore, since Q is compact, $(\Lambda(h_t, \mathbf{h}_\tau), \nu(\cdot | \mathbf{h}_\tau))_{\tau \in \mathbb{N}}$ admits $\mathbb{P}_{\Pi^{a^*}}$ a.s. a subsequence convergent to $(0, \nu^*)$ for some $\nu^* \in \Delta(Q(a^*))$. With this, the result follows from Lemma 3.

2) Suppose by contradiction that $a^* \notin BR^{Meu} \left(\left\{ p \in \Delta(Y)^A : \exists q \in Q, \forall a \in A, q_a \gg p_a \right\} \right)$ and that a^* is a Λ -limit action. Thus, since for every policy $\Pi \in A^{\mathcal{H}}$

$$\mathbb{P}_{\Pi}[\sup\{\mathbf{t} : \mathbf{a}_{\mathbf{t}} \neq a^*\} < \infty] \leq \sum_{t=0}^{\infty} \sum_{h_t \in \mathcal{H}_t} \mathbb{P}_{\Pi}[a^* = \Pi(h_\tau), \forall \tau \geq t | h_t] \mathbb{P}_{\Pi}[h_t],$$

there are a Λ -optimal policy $\tilde{\Pi} \in A^{\mathcal{H}}$, $t \in \mathbb{N}_0$, and $h_t \in \mathcal{H}_t$ with $\mathbb{P}_{\tilde{\Pi}}[h_t] > 0$ such that with positive probability $\tilde{\Pi}$ prescribes a^* after h_t in every future period. Define $\nu = \mu(\cdot | h_t)$, and notice that by Assumption 1 (i) $\text{supp } \nu = \text{supp } \mu = Q$. As the evolution of beliefs and misspecification concern under Π^{a^*} , i.e., the policy that plays a^* in every period, is the same as under $\tilde{\Pi}$ for every history where the agent continues to play a^* ,

$$\mathbb{P}_{\tilde{\Pi}}[a^* = \tilde{\Pi}(\mathbf{h}_\tau), \forall \tau > t | h_t] > 0 \implies \mathbb{P}_{\Pi^{a^*}}[a^* \in BR^{\Lambda(h_t, \mathbf{h}_\tau)}(\nu(\cdot | \mathbf{h}_\tau)), \forall \tau > t] > 0.$$

I now show that the latter equals zero, which establishes that a^* cannot be a Λ -limit action.

Since Y is a compact metric space, it is separable, and thus, by Varadarajan Theorem $\lim_{\tau \rightarrow \infty} p_{a^*}^{\mathbf{h}_\tau} = p_{a^*}^*$, $\mathbb{P}_{\Pi^{a^*}}$ -a.s. Then, by Lemmas 7 and 6, and equation (5),

$$\lim_{\tau \rightarrow \infty} \Lambda(h_t, \mathbf{h}_\tau) = \infty \quad \mathbb{P}_{\Pi^{a^*}}\text{-a.s.}$$

By Assumption 1 (i) for all $q, q' \in Q$ and $a \in A$, $q_a \sim q'_a$. So I obtain

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

Therefore, by Lemma 1 for all $a \in A$, $\mathbb{P}_{\Pi^{a^*}}$ -a.s. $\lim_{\tau \rightarrow \infty} \sup_{q \in Q} \min_{p_a \in \Delta(Y)} \int_Y u(a, y) dp_a + \frac{R(p_a || q_a)}{\Lambda(h_t, \mathbf{h}_\tau)} = \min_{y \in \cup_{q \in Q} \text{supp} q_a} u(a, y)$.

But since by Assumption 1 (i) for all $\tau \in \mathbb{N}$, $\mu(\cdot | \mathbf{h}_\tau) \subseteq Q$, $\mathbb{P}_{\Pi^{a^*}}$ -a.s.,

$$\lim_{\tau \rightarrow \infty} \mathbb{E}_{\mu(\cdot | \mathbf{h}_\tau)} \left[\min_{p_a \in \Delta(Y)} \int_Y u(a, y) dp_a + \frac{R(p_a || q_a)}{\Lambda(h_t, \mathbf{h}_\tau)} \right] = \min_{y \in \cup_{q \in Q} \text{supp} q_a} u(a, y) \quad \mathbb{P}_{\Pi^{a^*}}\text{-a.s.}$$

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

3) It follows from the more general Theorem 3. ■

Proof of Proposition 1. Observe that

$$R(p_a^* || q_a^\theta) = \ln \left(\frac{\sigma_2}{\sigma_1} \right) + \frac{\sigma_1^2 + (F(a) - \tau(a))^2}{2\sigma_2^2} - \frac{1}{2}, \quad (11)$$

and so $Q(a) = \{q^{F(a)}\}$, showing that there exists at most one Berk-Nash equilibrium. 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

$$(a - a')\theta = \mathbb{E}_{\varepsilon_{2,a}} [a(\theta + \varepsilon_{2,a})] - \mathbb{E}_{\varepsilon_{2,a'}} [a'(\theta + \varepsilon_{2,a'})] \geq C(a) - C(a').$$

By Proposition 1.4.2 in Dupuis and Ellis (2011), the condition for not switching from an action

a to an action a' with $a \geq a'$ in a c -robust equilibrium in which the belief is concentrated on θ is

$$\begin{aligned}
& (a - a') \theta + \frac{-c \log \mathbb{E}_{\varepsilon_{2,a}} \left[\exp \left(\frac{R(p_a^* || q_a^\theta) a \varepsilon_{2,a}}{-c} \right) \right]}{R(p_a^* || q_a^\theta)} + \frac{c \log \mathbb{E}_{\varepsilon_{2,a'}} \left[\exp \left(\frac{R(p_a^* || q_a^\theta) a' \varepsilon_{2,a'}}{-c} \right) \right]}{R(p_a^* || q_a^\theta)} \\
&= \frac{-c}{R(p_a^* || q_a^\theta)} \log \mathbb{E}_{\varepsilon_{2,a}} \left[\exp \left(-\frac{R(p_a^* || q_a^\theta) a \theta}{c} \right) \right] + \frac{-c}{R(p_a^* || q_a^\theta)} \log \mathbb{E}_{\varepsilon_{2,a}} \left[\exp \left(-\frac{R(p_a^* || q_a^\theta) a \varepsilon_{2,a}}{c} \right) \right] \\
&\quad + \frac{c}{R(p_a^* || q_a^\theta)} \log \mathbb{E}_{\varepsilon_{2,a'}} \left[\exp \left(-\frac{R(p_a^* || q_a^\theta) a' \theta}{c} \right) \right] + \frac{c}{R(p_a^* || q_a^\theta)} \log \mathbb{E}_{\varepsilon_{2,a'}} \left[\exp \left(-\frac{R(p_a^* || q_a^\theta) a' \varepsilon_{2,a'}}{c} \right) \right] \\
&= \frac{-c}{R(p_a^* || q_a^\theta)} \log \mathbb{E}_{\varepsilon_{2,a}} \left[\exp \left(-\frac{R(p_a^* || q_a^\theta) a \theta}{c} \right) \exp \left(-\frac{R(p_a^* || q_a^\theta) a \varepsilon_{2,a}}{c} \right) \right] \\
&\quad + \frac{c}{R(p_a^* || q_a^\theta)} \log \mathbb{E}_{\varepsilon_{2,a'}} \left[\exp \left(-\frac{R(p_a^* || q_a^\theta) a' \theta}{c} \right) \exp \left(-\frac{R(p_a^* || q_a^\theta) a' \varepsilon_{2,a'}}{c} \right) \right] \\
&= \frac{-c \log \mathbb{E}_{\varepsilon_{2,a}} \left[\exp \left(\frac{R(p_a^* || q_a^\theta) a (\theta + \varepsilon_{2,a})}{-c} \right) \right]}{R(p_a^* || q_a^\theta)} + \frac{c \log \mathbb{E}_{\varepsilon_{2,a'}} \left[\exp \left(-\frac{R(p_a^* || q_a^\theta) a' (\theta + \varepsilon_{2,a'})}{c} \right) \right]}{R(p_a^* || q_a^\theta)} \geq C(a) - C(a').
\end{aligned}$$

By the usual CARA-Gaussian formula, $a \geq a'$ implies that the LHS is lower in the second case, and I obtain the desired conclusion. \blacksquare

Lemma 8. For every $c \in \mathbb{R}_{++}$ the function $\alpha \mapsto \min_{q \in Q} \sum_{a \in A} \alpha(a) R(p_a^* || q_a) / c$ is continuous and the correspondence $Q(\cdot) : \Delta(A) \rightarrow 2^Q$ is upper hemicontinuous.

Proof of Proposition 2. Consider the following three-player game. The action sets are $A_1 = \Delta(A)$, $A_2 = \Delta(Q)$, $A_3 = \mathbb{R}_+$ with arbitrary elements denoted as α, ν, λ . The utility functions are

$$U_1(\alpha, \nu, \lambda) = \begin{cases} \sum_{a \in A} \alpha(a) \int_Q \min_{p_a \in \Delta(Y)} \left(\mathbb{E}_{p_a} [u(a, y)] + \frac{R(p_a || q_a)}{\lambda} \right) d\nu(q) & \lambda \neq 0 \\ \sum_{a \in A} \alpha(a) \int_Q \mathbb{E}_{q_a} [u(a, y)] d\nu(q) & \lambda = 0, \end{cases}$$

$U_2(\alpha, \nu, \lambda) = - \int_Q \sum_{a \in A} \alpha(a) R(p_a^* || q_a) d\nu(q)$, and $U_3(\alpha, \nu, \lambda) = - (\lambda - \min_{q \in Q} \sum_{a \in A} \alpha(a) R(p_a^* || q_a) / c)^2$.

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

$$\bar{\lambda} := \frac{\max_{\alpha \in \Delta(A)} \min_{q \in Q} \sum_{a \in A} \alpha(a) R(p_a^* || q_a)}{c} < \infty,$$

where the inequality holds by Assumption 1 (i). Therefore, since $\Delta(Q)$ is compact by the compactness of Q all the 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 3. Moreover, U_2 is trivially continuous in its first and third argument, and U_3 is continuous in its first and second argument by Lemma 8. 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

$$\lambda^* \in \operatorname{argmax}_{\lambda \in \mathbb{R}_+} - \left(\lambda - \min_{q \in Q} \sum_{a \in A} \alpha^*(a) R(p_a^* || q_a) / c \right)^2 \implies \lambda^* = \frac{\min_{q \in Q} \sum_{a \in A} \alpha^*(a) R(p_a^* || q_a)}{c},$$

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

and

$$\nu^* \in \operatorname{argmax}_{\nu \in \Delta(Q)} - \int_Q \sum_{a \in A} \alpha^*(a) R(p_a^* || q_a) d\nu(q) \implies \nu^* \in \Delta(Q(\alpha^*)).$$

Therefore, α^* is a mixed c -robust equilibrium sustained by ν^* and λ^* . ■

Proof of Theorem 3. I start observing that by Lemma 4, $\max_{q \in Q, t \in \mathbb{N}} -\ln \tilde{q}_{\mathbf{a}t}(\mathbf{y}_t) \leq K$, \mathbb{P}_{Π} -a.s. This will allow us to invoke Lemma 6.

Let $\Upsilon = \{\alpha - \alpha' : \alpha, \alpha' \in \Delta(A)\}$ and for all $\varepsilon \in \mathbb{R}_+$ and $\alpha' \in \Delta(A)$,

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

By Esponda et al., 2021a, Part 1a of the proof of Theorem 2,²² $M_{(\cdot)}(\cdot)$ is upper hemicontinuous.

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

$$F(\varepsilon, \alpha) = \left\{ \iota : \begin{array}{l} \exists \hat{\alpha}, \bar{\alpha} \in B_\varepsilon(\alpha), \lambda \in B_\varepsilon\left(\min_{q \in Q} \sum_{a \in A} \hat{\alpha}(a) R(p_a^* || q_a) / c\right), \nu \in \Delta(M_\varepsilon(\bar{\alpha})), \\ \iota \in \Delta(BR^\lambda(\nu)) - \alpha \end{array} \right\}$$

²²Observe that Assumption 1 (i-ii) implies their Assumption 2 (ii-iii), except for the fact that I do not require finite dimensionality of Y and Q . It is readily checked that since they are still assumed to be compact this does not create any issues in the proof of their Theorem 2.

and $\chi_\alpha = F(0, \alpha) + \alpha$.

Claim 5. F and χ_α are upper hemicontinuous and compact-valued.

Proof of the Claim. I show that F has a closed graph to conclude that it is upper hemicontinuous and compact-valued. Since Υ is compact, this is enough by, e.g., Proposition E.3 in Ok (2011). Let

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

be such that $\iota_n \in F(\varepsilon_n, \alpha_n)$ for all $n \in \mathbb{N}$ and convergent to $(\iota, \varepsilon, \alpha)$. Since A is finite, it is without loss of generality (possibly truncating some initial elements of the sequence) 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 $(\hat{\alpha}_n, \bar{\alpha}_n, \nu_n^{\hat{a}}, \lambda_n^{\hat{a}})_{n \in \mathbb{N}} \in (\Delta(A)^2 \times \Delta(Q) \times [0, 2K/c])^{\mathbb{N}}$ such that $\exists \hat{\alpha}_n, \bar{\alpha}_n \in B_{\varepsilon_n}(\alpha_n)$, $\nu_n^{\hat{a}} \in M_{\varepsilon_n}(\bar{\alpha}_n)$, $\lambda_n^{\hat{a}} \in B_\varepsilon(\min_{q \in Q} \sum_{a \in A} \hat{\alpha}_n(a) R(p_a^* || q_a) / c)$, and $\hat{a} \in BR^{\lambda_n^{\hat{a}}}(\nu_n^{\hat{a}})$. Since $\Delta(Q)$ and $[0, 2K/c]$ and $\Delta(A)$ are compact (by Theorem 15.11 in Aliprantis and Border, 2013) by restricting to a subsequence I can take $(\hat{\alpha}_n, \bar{\alpha}_n, \nu_n^{\hat{a}}, \lambda_n^{\hat{a}})_{n \in \mathbb{N}}$ to be convergent to some $(\hat{\alpha}, \bar{\alpha}, \nu^{\hat{a}}, \lambda^{\hat{a}}) \in \Delta(A)^2 \times \Delta(Q) \times [0, 2K/c]$. Since $M_{(\cdot)}(\cdot)$ is upper hemicontinuous $\nu^{\hat{a}} \in M_\varepsilon(\bar{\alpha})$. Since by Lemmas 8 and 3 $\alpha \mapsto \min_{q \in Q} \sum_{a \in A} \alpha(a) R(p_a^* || q_a) / c$ is continuous and $BR^{(\cdot)}(\cdot)$ is upper hemicontinuous,

$$\lambda^{\hat{a}} \in \min_{q \in Q} \sum_{a \in A} \hat{\alpha}(a) R(p_a^* || q_a) / c$$

and

$$\hat{a} \in BR^{\lambda^{\hat{a}}}(\nu^{\hat{a}}) \subseteq \left\{ BR^{\hat{\lambda}}(\hat{\nu}) : \hat{\nu} \in M_\varepsilon(B_\varepsilon(\alpha)), \hat{\lambda} \in \left\{ B_\varepsilon \left(\min_{q \in Q} \sum_{a \in A} \hat{\alpha}_n(a) R(p_a^* || q_a) / c \right) : \hat{\alpha} \in B_\varepsilon(\alpha) \right\} \right\}$$

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

Observe that $(\alpha_t)_{t \in \mathbb{N}}$ satisfies the following differential inclusion: for all $a \in A$, $t \in \mathbb{N}$, $h_t \in \mathcal{H}_t$, and $h_{t+1} \in \mathcal{H}_{t+1}$ such that $h_{t+1} \succ h_t$

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

Set $\tau_0 = 0$ and $\tau_t = \sum_{i=1}^t \frac{1}{i}$ for all $t \in \mathbb{N}$. The continuous-time interpolation of α_t is the

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

$$w(\tau_t + l) = \begin{cases} \alpha_t + l \frac{\alpha_{t+1} - \alpha_t}{\tau_{t+1} - \tau_t}, & \forall t \in \mathbb{N}, \forall l \in [0, \frac{1}{t+1}] \\ \alpha_1 & t = 0, \forall l \in [0, 1]. \end{cases} \quad (12)$$

I use the theory of stochastic approximation for differential inclusions (Benaïm et al., 2005 and Esponda et al., 2021a) to show that (12) can be approximated by a solution to

$$\dot{\underline{\alpha}}_t \in \chi_{\alpha_t} - \underline{\alpha}_t. \quad (13)$$

A solution over $[0, T]$, $T \in \mathbb{R}_{++}$, to the differential inclusion (13) with initial point $\hat{\alpha} \in \Delta(A)$ is a mapping $\underline{\alpha}_{(\cdot)} : [0, T] \rightarrow \Delta(A)$ that is absolutely continuous over compact intervals such that $\underline{\alpha}_0 = \hat{\alpha}$ and (13) is satisfied for almost every t . Let $S_{\hat{\alpha}}^T$ be the set of the solutions to (13) over $[0, T]$, $T \in \mathbb{R}_{++}$, with initial conditions $\hat{\alpha} \in \Delta(A)$. A solution to (13) exists by Claim 5 and Theorem 2.1.4 in Aubin and Cellina (2012), i.e., $S_{\hat{\alpha}}^T$ is nonempty for every $T \in \mathbb{R}_{++}$ and $\hat{\alpha} \in \Delta(A)$. Let $S^T = \cup_{\hat{\alpha} \in \Delta(A)} S_{\hat{\alpha}}^T$.

Observe that w is Lipschitz continuous of order 1 as for all $(h_t)_{t \in \mathbb{N}} \in \times_{t \in \mathbb{N}} \mathcal{H}_t$,

$$h_{t+1} \succ h_t \quad \forall t \in \mathbb{N} \implies \frac{\|\alpha_{t+1}(h_{t+1}) - \alpha_t(h_t)\|_{\infty}}{\tau_{t+1} - \tau_t} \leq \frac{1/(t+1)}{1/(t+1)} = 1 \quad \forall t \in \mathbb{N}. \quad (14)$$

Therefore w is absolutely continuous, as for all $n \in \mathbb{N}$ and $(l_i, \tilde{l}_i)_{i=1}^n$, $\sum_{i=1}^n \|w(l_i) - w(\tilde{l}_i)\|_{\infty} \leq \sum_{i=1}^n |l_i - \tilde{l}_i|$. Moreover, α_t is uniformly bounded because it takes values in $\Delta(A)$.

By Claim 5, χ_{α} satisfies Hypothesis 1.1 in Benaïm et al. (2005). Moreover, by Theorem 1 in Esponda et al. (2021a) and Lemma 6, \mathbb{P}_{Π} -almost surely, eventually $\mu(\cdot | \mathbf{h}_t) \in M_{\varepsilon}(\alpha_t(\mathbf{h}_t))$ and

$$\Lambda(\mathbf{h}_t) \in B_{\varepsilon} \left(\min_{q \in Q} \sum_{a \in A} \alpha_t(\mathbf{h}_t)(a) R(p_a^* | q_a) / c \right)$$

for all $\varepsilon \in \mathbb{R}_{++}^2$. Thus, there is a sequence $(\hat{\varepsilon}_t(\mathbf{h}_t))_{t \in \mathbb{N}} \in \mathbb{R}_{++}^{\mathbb{N}}$ converging to 0 with $\chi_{\alpha_t(\mathbf{h}_t)} - \alpha_t(\mathbf{h}_t) \in F(\hat{\varepsilon}_t(\mathbf{h}_t), \alpha_t(\mathbf{h}_t))$.

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

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

Claim 6. *Every limit point of the restrictions of $(G^t(w))_{t \in \mathbb{N}}$ on $[0, T]$ is in S^T .*

This argument borrows extensively from the proofs Theorem 4.2 in Benaim et al. (2005) and Theorem 2 in Esponda et al. (2021a). However they cannot be directly applied, because the interpolated process w I consider is not a perturbed solution in the sense of Benaim et al. (2005). Indeed, it may not be possible to find an α that *jointly* justifies a_t as a best reply to beliefs in $Q(\alpha)$ and the concern for misspecification $\min_{q \in Q} \sum_{a \in A} \alpha(a) R(p_a^* | q_a) / c$, as perturbations of the empirical frequency α_{t-1} in different directions may be needed for the concern and the belief. Nevertheless, the core of their arguments can be adapted by leveraging the upper hemicontinuity of F established above.

Proof of the Claim. Since w is uniformly continuous by equation (14), the family $(G^t(w))_{t \in \mathbb{N}}$ is equicontinuous, and thus it is relatively compact in the topology of uniform convergence over compact sets by the Arzela-Ascoli theorem (see Willard, 2012 Theorem 43.15 for the version with a noncompact domain). The topology of uniform convergence over compact sets is metrizable since $\Delta(A)$ is metrizable and \mathbb{R} is open (see Theorem 1.14b in Simon, 2020), and so there exists a limit point $z = \lim_{t_n} G^{t_n}(w)$. Define

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

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

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

Since $(v_n)_{n \in \mathbb{N}}$ is uniformly bounded, it is bounded in $L^2([0, T], \mathbb{R}^A, Leb)$. By the Banach-Alaoglu theorem (see Theorem 6.21 in Aliprantis and Border, 2013), (by restricting to a subsequence) I can take $(v_n)_{n \in \mathbb{N}}$ to be a weakly-convergent subsequence with limit $v^* \in L^2([0, T], \mathbb{R}^A, Leb)$. By Mazur's lemma (see Corollary V.3.14 in Dunford and Schwartz (1988)), there exist a function $N : \mathbb{N} \rightarrow \mathbb{N}$ and a sequence of positive weights $(\rho_n(n), \dots, \rho_{N(n)}(n))_{n \in \mathbb{N}}$ with $\sum_{i=n}^{N(n)} \rho_i(n) = 1$ for all $n \in \mathbb{N}$ such that if I define

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

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

Let $\tau \in [0, T]$ be such that $\lim_{n \rightarrow \infty} \bar{v}_n(\tau) = v^*(\tau)$. For every $t \in [0, T]$ and $n \in \mathbb{N}$, define

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

and

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

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

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

But then, by the upper hemicontinuity of F , for every $\varepsilon \in \mathbb{R}_{++}$ there exists N_ε such that for $n \geq N_\varepsilon$, $F(\gamma_n(t), w_n(t)) \subseteq B_\varepsilon(F(0, z(t)))$, where the latter set is closed and convex. But since $v_n(t) \in B_\varepsilon(F(0, z(t)))$, for all $n \geq N_\varepsilon$ also $\bar{v}_n(t) \in B_\varepsilon(F(0, z(t)))$. Therefore, $v^* \in F(0, z(\tau))$. Since the fact that v_n is weakly convergent to v^* implies by definition that $\lim_{n \rightarrow \infty} \int_0^T v_n(s) ds = \int_0^T v^*(s) ds$, $z \in S^T$. \square

Therefore, by (ii) \implies (i) of Theorem 4.1 in Benaim et al. (2005) (see Esponda et al., 2021b for the slightly corrected version used here)

$$\lim_{t \rightarrow \infty} \inf_{\underline{\alpha} \in S^T} \sup_{0 \leq s \leq T} \|w(t+s) - \underline{\alpha}_s\| = 0 \quad \mathbb{P}_\Pi\text{-a.s. for all } T \in \mathbb{N}. \quad (15)$$

With this, I can rule out convergence to nonequilibria. 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 $\chi_{(\cdot)}$ has a closed graph and maps into the compact $\Delta(A)$, 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 (13), $\underline{\alpha}(a)$ 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 $T^* := 4(D + \alpha^*(a)) / \alpha^*(a)$.

Next, I prove that $(\alpha_t(\mathbf{h}_t))_{t \in \mathbb{N}}$ does not converge to α^* on a sample path on which the convergence of equation (15) 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 $(\alpha_t(h_t))_{t \in \mathbb{N}}$ converges to α^* . Therefore, I can choose $\hat{T} \in \mathbb{N}$ such that on that sample path $\alpha_t(h_t) \in B_{D/2}(\alpha^*)$

for all $t \geq \hat{T}$ and

$$\inf_{\underline{\alpha} \in S^{T^*}} \sup_{0 \leq s \leq T^*} \|w(\hat{T} + s) - \underline{\alpha}_s\| \leq D/4. \quad (16)$$

Take any $\underline{\alpha} \in S^{T^*}$ with $\sup_{0 \leq s \leq T^*} \|w(\hat{T} + s) - \underline{\alpha}_s\| \leq D/2$. Since $w(\hat{T}) \in B_{D/2}(\alpha^*)$, $\underline{\alpha} \in S_{\bar{\alpha}}^{T^*}$ for some initial condition $\bar{\alpha} \in B_D(\alpha^*)$. But then by definition of T^* the differential inclusion leaves $B_D(\alpha^*)$ by time $T^* + \hat{T}$, and by (16), $(\alpha_t(\mathbf{h}_t))_{t \in \mathbb{N}}$ does not stay in $B_{D/2}(\alpha^*)$, a contradiction. ■

Proof of Corollary 1. Note that $\sum_{a \in A} \alpha(a) R(p_a^* \| q_a^\theta) = \sum_{a \in A} \alpha(a) \int \int \log \tilde{q}_a^\theta(y_U) dp_a^*(y_U | y_\pi) dp_a^*(y_\pi)$. By Assumption 3 (iii) for all $a \in A$, $\theta^* \in \cap_{y_\pi \in \text{supp} p_a^*} \min_{\theta \in \Theta} \int \log \tilde{q}_a^\theta(y_U) dp_a^*(y_U | y_\pi)$ and for every $\theta \in \Theta \setminus \{\theta^*\}$ there is a set $B \in \mathcal{B}(\mathbb{R})$ and an $\varepsilon > 0$ with $p_a^*(\{y \in Y : y_\pi \in B\}) > 0$ such that $\int \log \tilde{q}_a^\theta(y_U) dp_a^*(y_U | y_\pi) > \min_{\theta \in \Theta} \int \log \tilde{q}_a^\theta(y_U) dp_a^*(y_U | y_\pi)$ for all $y_\pi \in B$. Therefore for every $\alpha \in \Delta(A)$,

$$Q(\alpha) = \{q^{(\theta_0^*, \theta_{1\pi}^*, \theta_{1a}^*, \theta_2^*, \theta_3^*)}\}. \quad (17)$$

Moreover, since θ^* perfectly predicts the consequences under policy 0, $\min_{\theta \in \Theta} R(p_0^* \| q_0^\theta) = 0$. By Assumption 3 (i) and Lemma 3 in Battigalli et al. (2022), $BR^{Seu}(\Delta(Q(0))) = \{1\}$, and therefore 0 is not a c -robust equilibrium for any $c \in \mathbb{R}_{++}$. Since f_1 is strictly concave on \mathbb{R}_{++} , by equation (17) it follows that $\min_{\theta \in \Theta} R(p_1^* \| q_1^\theta) = R(p_1^* \| q_1^{\theta^*}) > 0$. By Assumption 3 (ii) and Lemma 1 there exists a sufficiently small \bar{c} such that for all $c \leq \bar{c}$, $BR^{\frac{\min_{\theta \in \Theta} R(p_1^* \| q_1^\theta)}{c}}(\delta_{\theta^*}) = \{0\}$ proving that there is no c -robust equilibrium if $c \leq \bar{c}$. That a mixed c -robust equilibrium exists follows from Proposition 2.

In particular, the maximal (resp. the minimal) equilibrium is defined as the α such that $\sum_{a \in A} \alpha(a) R(p_a^* \| q_a^{\theta^*}) / c$ is equal to the maximal (resp. minimal) misspecification concern λ such that $1 \in BR^\lambda(\delta_{\theta^*})$ (resp. $0 \in BR^\lambda(\delta_{\theta^*})$). Since, by Assumption 3 (iii), a larger $\theta_{1\pi}^* + \theta_{1a}^*$ makes action 0 more favorable, the comparative statics follows. ■

References

- Aitkin, M. (1991). “Posterior bayes factors”. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 53, pp. 111–128.
- Aliprantis, C. and K. Border (2013). *Infinite Dimensional Analysis: A Hitchhiker’s Guide*. Springer-Verlag.
- Anderson, E. W., L. P. Hansen, and T. J. Sargent (2003). “A quartet of semigroups for model specification, robustness, prices of risk, and model detection”. *Journal of the European Economic Association* 1, pp. 68–123.

- 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. (2024). “Robust Model Misspecification and Paradigm Shifts”. *arXiv preprint arXiv:2106.12727*.
- Barillas, F., L. P. Hansen, and T. J. Sargent (2009). “Doubts or variability?” *Journal of Economic Theory* 144, pp. 2388–2418.
- Battigalli, P. (1987). “Comportamento razionale ed equilibrio nei giochi e nelle situazioni sociali”. *Unpublished undergraduate dissertation, Bocconi University, Milano*.
- Battigalli, P., S. Cerreia-Vioglio, F. Maccheroni, M. Marinacci, and T. Sargent (2022). “A framework for the analysis of self-confirming policies”. *Theory and Decision*, pp. 1–58.
- Battigalli, P., A. Francetich, G. Lanzani, and M. Marinacci (2019). “Learning and self-confirming long-run biases”. *Journal of Economic Theory* 183, pp. 740–785.
- Benaim, M., J. Hofbauer, and S. Sorin (2005). “Stochastic approximations and differential inclusions”. *SIAM Journal on Control and Optimization* 44, pp. 328–348.
- Berk, R. H. (1966). “Limiting behavior of posterior distributions when the model is incorrect”. *The Annals of Mathematical Statistics*, pp. 51–58.
- Bohren, J. A. and D. N. Hauser (2021). “Learning with heterogeneous misspecified models: Characterization and robustness”. *Econometrica* 89, pp. 3025–3077.
- Cerreia-Vioglio, S., R. Corrao, and G. Lanzani (2024). “(Un-)Common Preferences, Ambiguity, and Coordination”. *mimeo*.
- Cerreia-Vioglio, S., L. P. Hansen, F. Maccheroni, and M. Marinacci (2024). “Making decisions under model misspecification”.
- Cerreia-Vioglio, S., G. Lanzani, and R. Corrao (2024). “Nonlinear Fixed Points and Stationarity: Economic Applications”. *Proceedings of the 25th ACM Conference on Economics and Computation*, pp. 400–400.
- Cerreia-Vioglio, S., F. Maccheroni, M. Marinacci, and L. Montrucchio (2013). “Ambiguity and robust statistics”. *Journal of Economic Theory* 148, pp. 974–1049.
- Cho, I.-K. and K. Kasa (2015). “Learning and model validation”. *The Review of Economic Studies* 82, pp. 45–82.
- Cho, I.-K. and J. Libgober (2023). “Learning Underspecified Models”. *arXiv preprint arXiv:2207.10140*.
- Cogley, T. and T. J. Sargent (2005). “The conquest of US inflation: Learning and robustness to model uncertainty”. *Review of Economic dynamics* 8, pp. 528–563.
- Dean, M. and P. Ortoleva (2019). “The empirical relationship between nonstandard economic behaviors”. *Proceedings of the National Academy of Sciences* 116, pp. 16262–16267.
- Dillenberger, D. and K. Rozen (2015). “History-dependent risk attitude”. *Journal of Economic Theory* 157, pp. 445–477.

- 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.
- Easley, D. and N. M. Kiefer (1988). “Controlling a stochastic process with unknown parameters”. *Econometrica*, pp. 1045–1064.
- Enke, B. and F. Zimmermann (2019). “Correlation neglect in belief formation”. *The Review of Economic Studies* 86, pp. 313–332.
- Epstein, L. G. and M. Schneider (2007). “Learning under ambiguity”. *The Review of Economic Studies* 74, pp. 1275–1303.
- Esponda, I. (2008). “Behavioral equilibrium in economies with adverse selection”. *American Economic Review* 98, pp. 1269–91.
- Esponda, I. and D. Pouzo (2016). “Berk–Nash equilibrium: A Framework for Modeling Agents with Misspecified Models”. *Econometrica* 84, pp. 1093–1130.
- Esponda, I., D. Pouzo, and Y. Yamamoto (2021a). “Asymptotic behavior of Bayesian learners with misspecified models”. *Journal of Economic Theory* 195, p. 105260.
- (2021b). “Corrigendum: Asymptotic behavior of Bayesian learners with misspecified models”. *Journal of Economic Theory* 195, p. 105260.
- Frick, M., R. Iijima, and Y. Ishii (2023). “Belief Convergence under Misspecified Learning: A Martingale Approach”. *Review of the Economic Studies* 90, pp. 781–814.
- Fudenberg, D. and G. Lanzani (2023). “Which misspecifications persist?” *Theoretical Economics* 18, pp. 1271–1315.
- Fudenberg, D., G. Lanzani, and P. Strack (2021). “Limit points of endogenous misspecified learning”. *Econometrica* 89, pp. 1065–1098.
- (2023). “Pathwise Concentration Bounds for Misspecified Bayesian Beliefs”. *Theoretical Economics* 18, pp. 1585–1622.
- (2024). “Selective Memory Equilibrium”. *Journal of Political Economy*, *Forthcoming*.
- Fudenberg, D. and D. K. Levine (1993). “Self-confirming equilibrium”. *Econometrica*, pp. 523–545.
- Gagnon-Bartsch, T., M. Rabin, and J. Schwartzstein (2023). “Channeled Attention and Stable Errors”.
- Giacomini, R., V. Skreta, and J. Turen (2020). “Heterogeneity, inattention, and bayesian updates”. *American Economic Journal: Macroeconomics* 12, pp. 282–309.
- Gilboa, I. and D. Schmeidler (1989). “Maxmin expected utility with non-unique prior”. *Journal of Mathematical Economics* 18, pp. 141–153.
- Hansen, L. P. and T. J. Sargent (2001). “Robust control and model uncertainty”. *American Economic Review* 91, pp. 60–66.
- (2007). “Recursive robust estimation and control without commitment”. *Journal of Economic Theory* 136, pp. 1–27.
- (2011). “Robustness”. *Robustness*. Princeton university press.

- Hansen, L. P. and T. J. Sargent (2022a). “Risk, Ambiguity, and Misspecification: Decision Theory, Robust Control, and Statistics”.
- (2022b). *Structured Ambiguity and Model Misspecification*.
- Hansen, L. P., T. J. Sargent, G. Turmuhambetova, and N. Williams (2006). “Robust control and model misspecification”. *Journal of Economic Theory* 128, pp. 45–90.
- He, K. and J. Libgober (2024). “Misspecified Learning and Evolutionary Stability”.
- 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”.
- 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.
- 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.
- 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.
- Sargent, T. J. (2008). “Evolution and intelligent design”. *American Economic Review* 98, pp. 5–37.
- Serfozo, R. (1982). “Convergence of Lebesgue integrals with varying measures”. *Sankhyā: The Indian Journal of Statistics, Series A*, pp. 380–402.
- Simon, J. (2020). *Continuous Functions*. John Wiley & Sons.
- Snowberg, E. and L. Yariv (2021). “Testing the waters: Behavior across participant pools”. *American Economic Review* 111, pp. 687–719.
- Sobel, J. (1984). “Non-linear prices and price-taking behavior”. *Journal of Economic Behavior & Organization* 5, pp. 387–396.
- Spiegler, R. (2020). “Can agents with causal misperceptions be systematically fooled?” *Journal of the European Economic Association* 18, pp. 583–617.
- Strzalecki, T. (2011). “Axiomatic foundations of multiplier preferences”. *Econometrica* 79, pp. 47–73.
- Tversky, A. and D. Kahneman (1971). “Belief in the law of small numbers.” *Psychological bulletin* 76, p. 105.
- Willard, S. (2012). *General topology*. Courier Corporation.