The Dynamics of Misspecification Fear*

Giacomo Lanzani

MIT

February 2022 (INCOMPLETE AND EXTREMELY PRELIMINARY DRAFT PLEASE DO NOT CIRCULATE)
Click here for the latest version.

Abstract

We consider a Decision Maker (henceforth DM) that posits a set of structural models describing the possible probability distributions over payoff relevant states. The DM has a probabilistic belief over this set, but they still fear that the true model is not in the support, and use a generalization of the multiplier preferences introduced by Hansen and Sargent (2001) to account for this concern. The agent uses Bayesian updating to adjust their belief about the models in light of the observed evidence. At the same time, the concern for misspecification is let to depend on the observed data. If there is a model that explains well the previous observations, the DM attenuates their concern for misspecification. We show how several (single-agent) versions of equilibrium concepts arise as the limit behavior, depending on the preferences of the DM and on whether the correct model is in the support.

*I am particularly indebted to Roberto Corrao, Drew Fudenberg, and Stephen Morris. I am also grateful to Abhijit Banerjee, Simone Cerreia-Vioglio, Joel Flynn, Massimo Marinacci, Philipp Strack, and Tomasz Strzalecki for helpful comments and suggestions. The financial support of Guido Cazzavillan Scholarship is gratefully acknowledged.

1 Introduction

The consideration of different probabilistic descriptions of reality and the use of the probability laws to adjust the relative weight assigned to each of these descriptions is the cornerstone of Bayesian rationality. However, even agents who correctly perform Bayesian between these probabilistic descriptions may fear that none of them is correct; they may be concerned to be misspecified.

Misspecification refers to when the Decision Maker (henceforth DM) does not assign a positive probability to the correct data generating process. It has been analyzed from two distinct perspectives, relying on an analytical separation between the concern for misspecification and the learning rule of the agent. On the one hand, several papers have studied the long-run implications of learning with misspecified beliefs, see, Esponda and Pouzo (2016), Esponda and Pouzo (2019), Frick, Iijima, and Ishii (2020) and Fudenberg, Lanzani, and Strack (2020). On the one other hand, the robust control literature in macroeconomics has proposed decision criteria for an agent who fears to be misspecified (see Hansen and Sargent (2001) and Hansen and Sargent (2020), respectively axiomatized by Strzalecki (2011) and Cerreia-Vioglio, Hansen, Maccheroni, and Marinacci (2020)). The main issue with these two criteria is that they are ill-suited for dynamic learning: In Hansen and Sargent (2020) there is no probabilistic belief over the set of conceived models, and because the DM uses a worst-case rule and therefore, except for the case of hard evidence against one of the models there is no way to update the decision criterion in the face of new evidence. Instead, the reference model is taken as given in the robust control model.

In this work, we reconcile these approaches and show how popular decision criteria such as the maxmin model, the robust control model, and subjective expected utility arise as the limit behavior of an agent concerned about misspecification and learning about the actual data generating process.

We consider a decision-maker that evaluates an act whose payoff depends on exogenous contingencies at each period. This evaluation is performed using a weighted average of robust control assessments, where each of these assessments takes as a reference measure a different structural model. We can now model learning with Bayesian updating of the probability measure on the structural models.

Importantly, we let the concern for misspecification be endogenous: the better the conceived models have explained the past evidence, the less the agent is concerned with misspecification. This is captured by considering agents who repeatedly perform likelihood ratio tests of their model and that adjust their concern for misspecification monotonically with respect to the resulting likelihood ratio statistics.

There will be two critical determinants for the long-run dynamics: whether the agent is correctly specified and the speed at which the agent adjusts her beliefs concern for misspecification. In particular, an agent who understands the central limit theorem applies a time-dependent normalization to the likelihood ratio statics that keeps it informative about the model's fitness over time. However, to capture the widespread evidence that many agents believe in the Law of Small Numbers (henceforth LSN) Tversky and Kahneman (1971) we also for agents that apply time normalization that treats the early period as a statistician would treat later periods. For completeness, we also consider the opposite case in which the agents are too lenient in evaluating their model.

The first result involves the particular case of a correctly specified agent. Both the lenient and the statistician types converge to a selfconfirming equilibrium: they play the best reply to a belief supported over the data generating process that it is observationally equivalent to the correct one given the chosen action. Instead, a believer in the LSN will become excessively concerned by the small imperfection of their model and will converge to a genuine robust control model with a positive concern for misspecification, even if there is no misspecification.

A more interesting taxonomy is obtained in the critical case of a misspecified agent. In that case, the actions of the lenient type converge to a Berk-Nash equilibrium, i.e., to an SEU best reply to beliefs supported on the models that are closer in terms of KL-divergence to the actual data generating process.

Instead, a statistician type will maintain a non-trivial amount of concern for misspecification. If their behavior converges, it converges to an average robust control best reply to the models that are closer in terms of KL-divergence to the actual data generating process. Notably, the behavior of such type is not guaranteed to converge. Indeed, it is possible that the behavior cycles between phases of different misspecification concerns. Loosely speaking, the agent can spend time playing an action whose consequences are explained highly by one of their structural models. The time spent playing this action lowers their concern for misspecification and eventually leads them to change behavior in favor of a more misspecification vulnerable action. Failure to explain the distribution of outcomes observed under this action leads to a return to the safe action. We apply this result to explain cyclical behavior for monetary policy and career concerns in organizations.

Finally, a misspecified believer in the LSN ends up overemphasizing the model's failures in explaining the data and converges to a maxmin equilibrium: it plays a maxmin best reply to the structural models that are absolutely continuous with respect to the true one. Summing up, the analysis provide a novel prediction that ambiguity aversion and the belief in the LSN should be positively correlated.

We also provide two novel axiomatic results: An axiomatization of the static average robust control criterion and a testable axiom for when the agent is of the lenient, statistician, or LSN believer type.

2 Static Decision Criterion

2.1 Static Average Robust Control

We start by describing the decision criterion used by the agent in our repeated decision problem, and we defer to Section 4 its axiomatization. We consider an agent who has a standard utility index $u : X \to \mathbb{R}$ over the set of certain outcomes X that captures their preference when the subjective uncertainty is resolved. However, the realized outcome is stochastic, and depends on underlying state $s \in S$. The agent deals with this uncertainty in a quasi-Bayesian way. They postulate a set $Q \subseteq \Delta(S)$ of possible *structural models,* i.e., Borel probability measures over states $q \in \Delta(S)$, and they have a belief $\mu \in \Delta(Q)$ that describes the relative likelihood assigned to these models.[1] For example, if the agent is a central bank they may conceive both a Keynesian Samuelson-Solow model where the monetary policy affects the unemployment rate, or a new classical Lucas-Sargent model with no effect of systematic inflation on unemployment. However, the agent is concerned with the possibility that none of these probability distributions is the correct description of the data generating process, but only a valuable approximation. Therefore, in the spirit of the robustness criterion advocated by Hansen and Sargent (2001), they penalize acts who perform badly under alternative distributions.

In particular, as in Hansen and Sargent (2001), the trade-off between robustness of the decision and the performance under the structural model $q \in \Delta(S)$ is governed by a parameter λ, and the concern about the performance under an alternative model $p \in \Delta(S)$ decreases in its Kullback-Leibler divergence from the structural model, given by

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

[1]We endow $\Delta(S)$ with the Prokhorov metric and the corresponding Borel sigma algebra.

where $q \gg p$ means that p is absolutely continuous with respect to q. With this, an agent with Bernoulli utility index $u : X \to \mathbb{R}$, belief $\mu \in \Delta(Q)$ over the set of structural models, and concern for misspecification $\lambda \in \mathbb{R}_{++}$ evaluates the act $f : S \to X$ accordingly to the *average robust control* criterion:

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

Several well-known decision criteria are obtained as special (or limit) cases of the average robust control representation. The limit for $\lambda \to 0$ corresponds to a (classical) subjective expected utility maximizer (Cerreia-Vioglio, Maccheroni, Marinacci, and Montrucchio, 2013). The limit for $\lambda \to \infty$ is the widely used multiple prior model axiomatized by Gilboa, Schmeidler, et al. (1989) where the set of priors is $C = \{p \in \Delta(S) : \exists q \in Q, q \gg p\}$. The standard multiplier preference introduced by Hansen and Sargent (2001) and axiomatized by Strzalecki (2011) is the particular case in which μ is a Dirac measure.

The representation clearly shows the existence of two sources of uncertainty. At a first level, given a probabilistic model q, the uncertainty about the exact specification of the model is captured by minimizing the value of the act f with respect to probabilities that are not too far away from q. At a higher level, the agent is also uncertain about the identity of the best structural model, and posits a probability μ over them.

3 Dynamic Environment

We are going to consider a set Ω of one period consequences observed by the agent. For every set Ω, we let $\Omega^t = \prod_{\tau=1}^{t} \Omega$ and $\Omega^\infty = \prod_{\tau=1}^{\infty} \Omega$.[2] We endow the space Ω^∞ with the Borel σ-algebra, $\mathcal{B}(\Omega^\infty)$, corresponding to the

[2]Unless otherwise stated, it is understood that t is an element of $\mathbb{N}$, the set of natural numbers. We use the terms "time" and "period" interchangeably to refer to t.

product topology on Ω^∞; this is the same as the σ-algebra generated by the elementary cylinders $\{\omega_1\} \times \cdots \times \{\omega_t\} \times \Omega^\infty$ (see, e.g., Proposition 1.3 in (Folland, 1999)). We denote by $\omega^t = (\omega_1, ..., \omega_t) \in \Omega^t$ both the histories and the elementary cylinders that they identify through the following map:

$$(\omega_1, ..., \omega_t) \mapsto \{\omega_1\} \times \cdots \times \{\omega_t\} \times \Omega^\infty \ .$$

We denote by $\omega^\infty = (\omega_1, ..., \omega_t, ...)$ a generic element of Ω^∞. Each one period consequence ω_i corresponds to a subset $E_{\omega_i} \subseteq S$.

There is a set of action A corresponding to the alternatives available to the decision maker at each period $t \in \{1, 2, 3, \ldots\}$. Each action $a \in A$ induces an objective probability distribution $p_a^* \in \Delta(\Omega)$ over the set of possible outcomes Ω. Moreover, the action, paired with the realized outcome, determines the flow payoff of the agent via the utility function $u : A \times \Omega \to \mathbb{R}$.

An history is a finite vector of past choice and consequences. In particular, $\mathcal{H}_t = A^t \times \Omega^t$ and $\mathcal{H} = \bigcup_{t=0}^\infty \mathcal{H}_t$.

Subjective Beliefs of the Agent The agent correctly believes that the map from actions to probability distributions over outcomes is fixed and depends only on their current action, but they are uncertain about the distribution each action induces. Let $P =_{a\in A} \Delta(\Omega)$ be the space of all action-dependent outcome distributions, and let $p_a \in \Delta(\Omega)$ denote the a-th component of $p \in P$. We endow P with the sup-norm topology, and denote by $B_\varepsilon(p)$ the ball of radius ε around $p \in P$.[3]

The agent's uncertainty is captured by a compact set of parameters $\Theta \subseteq \mathbb{R}^m$, $m \in \mathbb{N}$ and a prior belief $\mu_0 \in \Delta(\Theta)$, where $\Delta(\Theta)$ denotes the metric space of Borel probability measures on Θ endowed with the topology of weak convergence of measures. Each parameter $\theta \in \Theta$ is associated with a distribution $p^\theta \in P$

[3]For every finite dimensional vector v, we let $||v|| = \max_i v_i$ denote the supremum norm.

Definition 1. The agent is *correctly specified* if there exists $\theta \in \Theta$, with $p^{\theta} = p^{*}$, i.e. the objective distribution is conceivable.

Updating Subjective Beliefs We assume throughout that the agent updates their beliefs using Bayes rule. Denote by $\mu_t(\cdot \mid (a^t, \omega^t))$ the subjective belief the agent obtains using Bayes rule after action sequence $a^t = (a_s)_{s=1}^{t}$ and outcome sequence $\omega^t = (\omega_s)_{s=1}^{t}$,

$$\mu_t(C \mid (a^t, \omega^t)) = \frac{\int_{p \in C} \prod_{\tau=1}^{t} p_{a_\tau}(\omega_\tau) d\mu_0(p)}{\int_{p \in P} \prod_{\tau=1}^{t} p_{a_\tau}(\omega_\tau) d\mu_0(p)}. \qquad \text{(Bayes Rule)}$$

Since the agent's prior has support Θ, their posterior belief does as well.

Behavior of the Agent A (pure) policy $\pi : \bigcup_{t=0}^{\infty} A^t \times \Omega^t \to A$ specifies an action for every history. Throughout, we let $a_{t+1} = \pi(a^t, \omega^t)$ denote the action taken in period t. The objective action-contingent probability distribution p^* and a policy π induce a probability measure $\mathbb{P}_\pi$ on $(a_\tau, \omega_\tau)_{\tau=1}^{\infty}$.

For every $\lambda \in \mathbb{R}_{++}$ let $BR^{\lambda}(\nu)$ denotes the set of average robust control best replies to ν, i.e.,

$$BR^{\lambda}(\nu) = \operatorname*{argmax}_{a \in A} \mathbb{E}_{\nu} \left[\min_{p \in \Delta(\Omega)} \mathbb{E}_{p_a} \left[u(a, \omega) \right] \mathrm{d}p + \frac{1}{\lambda} R(p||q) \right].$$

Let also

$$BR^{0}(\nu) = \operatorname*{argmax}_{a \in A} \mathbb{E}_{\nu} \left[\mathbb{E}_{p_a} \left[u(a, \omega) \right] \right]$$

3.1 Equilibrium Concepts

Definition 2. An action is a selfconfirming equilibrium if there exists $\nu \in \Delta(\Theta)$ with su$pp\nu \subseteq \{p \in \Theta : p_a = p^*\}$ and $a \in BR^0(\nu)$.

3.2 Adjusting the concern for misspecification

The decision-maker is assumed to choose accordingly to

$$V(f|h_t) = \mathbb{E}_{\mu(\cdot|h_t)}\left[\min_{p\in\Delta(S)}\int_S\left[u(f)+\frac{R(p||q)}{\lambda_t(h_t,\mu)}\right]\mathrm{d}p\right]$$

where $\mu(\cdot|h_t)$ is computed using Bayesian updating, $\lambda_t(h_t,\mu) = \bar{G}_t(h_t,\Theta)$.We consider

$$\bar{G}_t^{\alpha_t}(h_t,\Theta) = \left(\log\left(\frac{\mathbb{P}_{p(h_t)}(h_t)}{\max_{\theta\in\Theta}\mathbb{P}_\theta(h_t)}\right)\right)/\alpha_t = \frac{\text{Likelihood Ratio Test Statistics(LRT)}}{2\alpha_t}.$$

Lemma 1. *If* $q \in \arg\max_{\theta\in\Theta}\mathbb{P}(p(h_t)|\theta))$, *we have*

$$\bar{G}_t^t(h_t,\Theta) = d_{KL}(p(h_t),q).$$

A policy π is α_t-optimal if for all h_t, $\pi(h_t) \in BR^{\bar{G}_t^{\alpha_t}(h_t,\Theta)}(\nu)$.

3.2.1 Full Support Agent

Proposition 1. *Let* Y *be finite and* $\Theta = \Delta(Y)^A$. *For every* $\alpha_t \in \mathbb{R}_{++}^{\mathbb{N}}$ *if*

$$\mathbb{P}_\pi[\lim a_t = a] > 0$$

then a *is a self-confirming equilibrium.*

3.2.2 General Case

Theorem 1. *Suppose that* $\text{supp}\mu$ *is finite,* $\{q\} = \arg\min_{p\in Q} d_{KL}(p^*,p)$

1. *If* $t = o(\alpha_t)$, *then* $\lim_{t\to\infty}\mathbb{P}\left(\mathbf{f}_t \in BR^{SEU}(\delta_{p^*})\right) = 1$.
2. *If* $\alpha_t = o(t)$, *and for all* $q,p \in \text{supp}\mu$, $q \sim p$ *then* $\lim_{t\to\infty}\mathbb{P}\left(\mathbf{f}_t \in BR^{GS}(\delta_{p^*})\right) = 1$.

3. *If* $\alpha_t = ct$ *for some* $c \in \mathbb{R}$ *and* $\{f\} = BR^{\lambda}(\delta_{p^*})$ *for some* $\lambda \in \mathbb{R}_+$*, then almost surely* $f \in \limsup_{t\to\infty} f_t$

$$\lim_{t\to\infty} \mathbb{P}\left(\mathbf{f}_t \in BR^{d_{KL}(p^*,q)}(\delta_{p^*})\right) = 1.$$

4 Representation

We aim to derive the average robust control representation from axioms over the binary relation $\succsim$.

4.1 Notation and Preliminaries

Consider a nonempty set of states S endowed with a separable σ-algebra of events Σ such that (S, Σ) is a standard Borel measurable space. The DM envisions the set of simple acts on (S, Σ), i.e., the Σ-measurable maps from states into a set of outcomes X with a finite range. The set of those acts is denoted as $\mathcal{F}$. We assume that X is a convex subset of a finite-dimensional vector space. Given any $x \in X$, $x \in \mathcal{F}$ is the act that delivers x in every state, and in this way, we identify X as the subset of constant acts in $\mathcal{F}$. If $f, g \in \mathcal{F}$, and $E \in \Sigma$, we denote as gEf the simple act that yields $g(s)$ if $s \in E$ and $f(s)$ if $s \notin E$. Since X is convex, for every $f, g \in \mathcal{F}$, and $\alpha \in (0,1)$, we denote as $\alpha f + (1-\alpha) g \in \mathcal{F}$ the simple act that pays $\alpha f(s) + (1-\alpha) g(s)$ for all $s \in S$. We denote as $\Delta(S)$ the space of all probability measures over (S, Σ), and we endow it with the topology of weak convergence of measures and the Prokhorov metric, and we let Σ_Δ be the corresponding Borel σ-algebra. We also denote as $\Delta(\Delta(S))$ the set of probability measures over $(\Delta(S), \Sigma_\Delta)$. For a set $C \subseteq \Delta(S)$ we let Σ_Δ^C be the restriction of Σ_Δ on C.

We model the decision maker's preference with a binary relation $\succsim$ on $\mathcal{F}$. As usual $\succ$ and $\sim$ denote the asymmetric and symmetric parts of $\succsim$. An event E is null if $fEh \sim gEh$ for every $f, g, h \in \mathcal{F}$. If $f \in \mathcal{F}$, an element $x_f \in X$ is a certain equivalent of f if $f \sim x_f$.

4.2 Decision Criterion

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

Definition 3. A tuple (u, Q, μ, λ) is an average robust control representation of the preference relation $\succsim$ if and only if $u : X \to \mathbb{R}$ is a nonconstant affine function, $\mu \in \Delta(Q)$, $Q \subseteq \Delta(S)$ is a nonempty set and

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

4.3 Static Axioms

The first minimal requirement for the preference relation is to be complete and transitive.

Axiom 1 (Weak Order)**.** *If $f, g, h \in \mathcal{F}$, (i) either $f \succsim g$ or $g \succsim f$, and (i) $f \succsim g$ and $g \succsim h$ imply $f \succsim h$.*

Next, we require the Weak Certainty Independence Axiom introduced by (Maccheroni, Marinacci, and Rustichini, 2006). It requires that although the agent may perceive some advantage in hedging, this cannot come from mixing with different constants using the same weights.

Axiom 2 (Weak Certainty Independence)**.** *If $f, g \in \mathcal{F}$, $x, y \in X$, and $\alpha \in (0, 1)$,*

$$\alpha f + (1-\alpha)\,x \succsim \alpha g + (1-\alpha)\,x \Rightarrow \alpha f + (1-\alpha)\,y \succsim \alpha g + (1-\alpha)\,y.$$

We also impose a standard technical continuity assumption.

Axiom 3 (Continuity)**.** *If $f, g, h \in \mathcal{F}$ the sets $\{\alpha \in [0,1] : \alpha f + (1-\alpha)\,g \succsim h\}$ and $\{\alpha \in [0,1] : h \succsim \alpha f + (1-\alpha)\,g\}$.*

The next Monotonicity assumption requires that the preference over act is minimally consistent with the preferences over the outcomes they induce, and that the utility of an outcome is not state dependent.

Axiom 4 (Monotonicity)**.** *If $f, g \in \mathcal{F}$ and $f(s) \succsim g(s)$ for all $s \in S$, then $f(s) \succsim g(s)$.*

The next axiom takes a stance on the attitudes about misspecification of the agent, leading to an aversion for the act that performs well with respect to a postulated model but poorly with respect to its perturbation.

Axiom 5 (Uncertainty Aversion)**.** *If $f, g \in \mathcal{F}$ and $\alpha \in (0,1)$*

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

We also need to assume that the problem is not trivial.

Axiom 6 (Nondegeneracy)**.** *$f \succ g$ for some $f, g \in \mathcal{F}$.*

The previous axioms where used to characterize the class of variational preferences by (Maccheroni, Marinacci, and Rustichini, 2006). We say that an event $E \subseteq S$ satisfies the sure-thing principle if, for all $f, g, h, h' \in \mathcal{F}$, the following conditions are satisfied

1. If $fEh \succsim gEh$, then $fEh' \succsim gEh'$.
2. If $hEf \succsim hEg$ then $h'Ef \succsim h'Eg$.

We denote as Σ_{st} the set of events that satisfies the sure-thing principle. For every nonnull $E \in \Sigma_{st}$, the conditional preference relation $\succsim_E$ is defined by $f \succsim_E g$ if $fEh \succsim gEh$ for some $h \in \mathcal{F}$.

Axiom 7 (Intramodel Sure-Thing Principle)**.** *For all non-null $E \in \Sigma_{st}$, $B \in \Sigma$ and $f, g, h, h' \in \mathcal{F}$*

$$fBh \succsim_E gBh \implies fBh' \succsim_E gBh'.$$

We also denote as $\mathcal{F}_{\Sigma_{st}}$ the set of acts that are measurable with respect to the partition Σ_{st}.

Axiom 8 (Weak Montone Continuity)**.** *If $f, g \in \mathcal{F}$, $x \in X$, $(E_i)_{i\in\mathbb{N}} \in \Sigma$ with $E_1 \supseteq E_2 \supseteq ...$ and $\cap_{n\geq 1} E_n = \emptyset$, then $f \succ g$ implies that there exists $n_0 \in \mathbb{N}$ such that $xE_{n_0}f \succ g$.*

An important condition for the collection of models Q is identifiability. Loosely speaking, it requires that there is a way to partition the state space that identifies the probabilistic model; Each of the probabilistic models assigns probability one to its corresponding element of the partition.

Definition 4. A nonempty set $Q \subseteq \Delta$ is identifiable if there exists a measurable function $k : S \to Q$ such that

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

An average robust control representation (u, Q, μ, λ) is *identifiable* if the set Q is identifiable and μ is nonatomic. In that case we say that k identifies Q.

To see this more concretely, observes that in a structured average robust control representation $k\left(\omega, \pi\right) = \pi$ identifies Q.

Theorem 2. *Suppose that every nonnull $E \in \Sigma_s$ contains at least three disjoint non-null events, then $\succsim$ admits an identifiable average robust control representation if and only if it satisfies Weak Order, Weak Certainty Independence, Continuity, Monotonicity, Uncertainty Aversion, Nondegeneracy, and the Intramodel Sure-Thing Principle.*

Proof of Theorem 2. (Only if) That $\succsim$ satisfies Weak Order, Weak Certainty Independence, Continuity, Monotonicity, Uncertainty Aversion, and Nondegeneracy follows by Lemma 2. It satisfies Intramodel Sure-Thing Principle by Lemma 16 in (Denti and Pomatto, 2020) and Proposition 7.

(If) It follows by Proposition 4 together with (Savage, 1972). ■

4.4 Relative Concern for Misspecification

A key notion to understand how the concern for misspecification evolves over time is a notion of being more misspecification concerned.

Definition 5. Given two preferences $\succsim_1$ and $\succsim_2$ on $\mathcal{F}$, we say that $\succsim_1$ is more concerned with misspecification than $\succsim_2$ if, for each $f \in \mathcal{F}$ and each $x \in X$, $f \succsim_1 x$ implies $f \succsim_2 x$.

An important benchmark is the case in which the decision maker may assign a different probability to each model being the best explanation of the data generating process but within each of the model has the same concern for not being the exact description of the world.

Axiom 9 (Model Independent Concern for Misspecification)**.** *For every* $E, E' \in \Sigma_{st}$, $x, y \in X$ *with* $x \succ y$ *and* $E \subset B, E' \subset B'$ *such that*

$$xEy \sim \frac{(x+y)}{2}By \text{ and } xE'y \sim \frac{(x+y)}{2}B'y$$

implies

$$\left\{z : y \succ z, zEy \succ \frac{(z+y)}{2}By\right\} = \left\{z : y \succ z, zE'y \succ \frac{(z+y)}{2}B'y\right\}.$$

Proposition 2. *Let* $\succsim$ *admit an average robust control representation. Then* $\succsim$ *satisfies Model Independent Concern for Misspecification if and only if there exists* λ^* *with* $\lambda_Q = \lambda^*$.

4.5 Dynamic Axioms

To describe how preference evolve in face of information, we will need to consider a family of preferences $\left(\succsim^H\right)_{H\in\mathcal{H}}$ indexed by the realized history.

Axiom 10 (Constant Preference Invariance)**.** *For any* $x, y \in X$, *and* $H \in \mathcal{H}$, $x \succsim^H y \Leftrightarrow x \succsim^\emptyset y$.

Axiom 11 (Dynamic Consistency over Structural Models)**.** *Suppose that for every nonnull* $E \in \mathcal{H}_1$ *and for every* $f, g \in \mathcal{F}$ *that are measurable with respect to* Σ_{st}

$$f \succsim^{\emptyset}_{E} g \Longleftrightarrow fBg \succsim^{E} g.$$

Axiom 12 (Q-Likelihood Ratio Agents)**.** *For every* $t \in \mathbb{N}$*,* $H, H' \in \mathcal{H}_t$*, and* $E \in \Sigma_s$*,* $x \in X$*, and* $f \in \mathcal{F}$*, if then*

$$LRT_Q(H) \geq LRT_Q(H') \text{ and } f \succsim_H x \implies f \succsim_{H'} x.$$

Proposition 3. *Let* $\left(\succsim^H\right)_{H \in \mathcal{H}}$ *and* $Q \subseteq \Delta(S)$ *be such that:*

1. *each* $\succsim^H$ *satisfies Weak Order, Weak Certainty Independence, Continuity, Monotonicity, Uncertainty Aversion, Nondegeneracy, and the Intramodel Sure-Thing Principle,*

2. *for each* $\succsim^H$*,* Q *is the smallest closed set* E *such* $S \times (\Delta(S) \setminus E)$*,*

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

Then each $\succsim^H$ *admits an average robust control representation* (u, Q, μ_H, λ_H) *with* λ_H *increasing in* LRT_Q*.*

5 Appendix

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

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

A niveloid is normalized if $I(k1_S) = k$ for all $k \in \mathbb{R}$ such that $k1_S \in \Phi$.

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

Lemma 2. *If $\succsim$ admits an average robust control representation (u, Q, μ, λ) then it satisfies Weak Order, Weak Certainty Independence, Continuity, Monotonicity, Uncertainty Aversion, and Nondegeneracy, and therefore it admits a representation of the form*

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

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

Proof By Theorem 3 and Lemma 28 in (Maccheroni, Marinacci, and Rustichini, 2006), for every $q \in \Delta(S)$, there exists a normalized niveloid $I_q : B_0(\Sigma, u(X))$ and a nonconstant affine function v_q such that

$$\min_{p \in \Delta(S)} \int_S \left[u(f) + \frac{1}{\lambda} R(p||q) \right] \mathrm{d}p = I_q(v_q(f))$$

and v_q can be assumed to be such that $0 \in \mathrm{int}(v_q(X))$. Let $q, q' \in \Delta(S)$, and $x \in X$. We have that

$$\begin{aligned} v_q(x) &= I_q(v_q(x)) = \min_{p \in \Delta(S)} \int_S \left[u(x) + \frac{1}{\lambda} R(p||q) \right] \mathrm{d}p = u(x) \\ &= \min_{p \in \Delta(S)} \int_S \left[u(x) + \frac{1}{\lambda} R(p||q') \right] \mathrm{d}p = I_{q'}(v_{q'}(x)) = v_{q'}(x) \end{aligned}$$

showing that $v_q = v_{q'} : = v$. By Lemma 25 in (Maccheroni, Marinacci, and Rustichini, 2006), I is monotone and translation invariant. Let $\mu \in$

$\Delta(\Delta(S))$. For every $\varphi \in B(\Sigma, v(X))$, define

$$\hat{I}(\varphi) = \int_{\Delta(\Delta(S))} I_q(\varphi)\, \mathrm{d}\mu(q).$$

By the monotonicity of the integral, $\hat{I}$ is monotone. Let $\varphi \in B_0(\Sigma, K)$, $k \in v(X)$, and $\alpha \in (0,1)$. We have

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

where the second equality follows from the translation invariance of each I_q. But then, notice that

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

where $\hat{I}$ is monotone and translation invariant. Therefore, the statement follows from Lemmas 25 and 28 and Theorem 3 of (Maccheroni, Marinacci, and Rustichini, 2006). By the first part of the lemma we have

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

and therefore by the uniqueness up to a positive affine transformation of $\hat{u}$ guaranteed by Theorem 3 of (Maccheroni, Marinacci, and Rustichini, 2006) and the fact that each two affine function that represent $\succsim$ on X are positive affine transformations of each other (Herstein and Milnor, 1953), we can choose $u = \hat{u}$. Finally, that the unique $\hat{c}$ identified by the choice of $\hat{u}$

has $\hat{c}^{-1}(0) = \mathbb{E}_\mu[p]$ follows by Lemma 32 of (Maccheroni, Marinacci, and Rustichini, 2006). ■

We then show that if in every state the outcome paid by act f is strictly preferred to the one paid by act g then f is strictly preferred to g.

Lemma 3. *If $\succsim$ satisfies Weak Order, Weak Certainty Independence, Continuity, Monotonicity, Uncertainty Aversion and Nondegeneracy, then if $f(s) \succ g(s)$ for all $s \in S$, then $f \succ g$.*

Proof If $\succsim$ satisfies Weak Order, Weak Certainty Independence, Continuity, Monotonicity, Uncertainty Aversion and Nondegeneracy, then by Theorem 3 of (Maccheroni, Marinacci, and Rustichini, 2006), it admits a variational representation:

$$f \succsim g \Longleftrightarrow \min_{q\in\Delta}\left(\int u(f)\,dq + c(q)\right) \geq \min_{q\in\Delta}\left(\int u(g)\,dq + c(q)\right)$$

for some nonconstant affine $u : X \to \mathbb{R}$ and a grounded, convex, lower semicontinuous function $c : \Delta(S) \to [0,\infty]$. Then the result follows immediately from the representation. ■

Lemma 4. *If $\succsim$ admits an average robust control representation then if $f(s) \succ g(s)$ for all $s \in S$, then $f \succ g$.*

Proof It follows immediately from Lemmas 2 and 3. ■

Lemma 5. *Let $E \in \Sigma_{st}$ be nonnull, and let $\succsim$ satisfy Weak Order, Weak Certainty Independence, Continuity, Monotonicity, Uncertainty Aversion, and Weak Monotone Continuity, then $\succsim_A$ satisfies Weak Order, Weak Certainty Independence, Continuity, Monotonicity, Uncertainty Aversion, Nondegeneracy, and Weak Monotone Continuity.*

Proof Let $f, g, h \in \mathcal{F}$. By Completeness of $\succsim$ at least one between $fEh \succsim gEh$ and $gEh \succsim fEh$ holds. Therefore, by definition of $\succsim_E$ at least one between $f \succsim_E g$ and $g \succsim_E f$ holds. Let $f, f', f'', h \in \mathcal{F}$, with $f \succsim_E f'$ and $f' \succsim_E f''$. Since $f \succsim_E f'$, $fEh' \succsim f'Eh'$ and $f'Eh'' \succsim f''Eh''$ for some $h', h'' \in \mathcal{F}$. Since $E \in \Sigma_{st}$, $fEh'' \succsim f'Eh''$. By Transitivity of $\succsim$, $fEh'' \succsim f''Eh''$, and so by definition of $\succsim_E$, $f \succsim_E f''$.

Let $f, g \in \mathcal{F}$, $x, y \in X$, and $\alpha \in (0,1)$, be such that $\alpha f + (1-\alpha) x \succsim_E \alpha g + (1-\alpha) x$. Since $E \in \Sigma_{st}$ we have $(\alpha f + (1-\alpha) x) Ex \succsim (\alpha g + (1-\alpha) x) Ex$. By Weak Certainty Independence of $\succsim$ we get $(\alpha f + (1-\alpha) y) E (\alpha x + (1-\alpha) y) \succsim (\alpha g + (1-\alpha) y) E (\alpha x + (1-\alpha) y)$. But then by definition of $\succsim_E$, we have $\alpha f + (1-\alpha) y \succsim_E \alpha g + (1-\alpha) y$, proving that $\succsim_E$ satisfies Weak Certainty Independence.

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

$$\{\alpha \in [0,1] : \alpha f + (1-\alpha) g \succsim_E h\} = \{\alpha \in [0,1] : (\alpha f + (1-\alpha) g) Eh' \succsim hEh'\}$$

and

$$\{\alpha \in [0,1] : h \succsim_E \alpha f + (1-\alpha) g\} = \{\alpha \in [0,1] : hEh' \succsim (\alpha f + (1-\alpha) g) Eh'\}$$

where the sets on the RHS' are closed by Continuity of $\succsim$.

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

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

Since E is nonnull, there exist $f, g, h \in \mathcal{F}$ such that $fEh \succ gEh$. But then, by definition of $\succsim_E$, $f \succ_E g$ and $\succsim_E$ satisfies Nondegeneracy.

Let $f, g \in \mathcal{F}$, $x \in X$, $(E_i)_{i \in \mathbb{N}} \in \Sigma$ with $E_1 \supseteq E_2 \supseteq ...$ and $\cap_{n \geq 1} E_n = \emptyset$, and $f \succ_E g$. Then $(E_i')_{i \in \mathbb{N}}$ where $E_i' = E_i \cap E$ is such that $E_1' \supseteq E_2' \supseteq ...$ and $\cap_{n \geq 1} E_n' = \emptyset$. Then $(fE_i'x) Eh = (fEh) E_i'x$ and by Weak Monotone

Continuity and the fact that $fEh \succ gEh$ there exists $n_0 \in \mathbb{N}$ such that $\left(fE'_{n_0}x\right)Eh \succsim gEh$. But notice that $\left(fE_{n_0}x\right)Eh = \left(fE'_{n_0}x\right)Eh \succsim gEh$ and therefore $fE_{n_0}x \succsim_E g$. ■

Proposition 4. *Let $E \in \Sigma_{st}$ contain at least three disjoint non-null events, and let $\succsim$ satisfy Weak Order, Weak Certainty Independence, Continuity, Monotonicity, Uncertainty Aversion, and Weak Monotone Continuity, and the Intramodel Sure-Thing Principle then*

$$f \succsim_E g \Longleftrightarrow \min_{q\in\Delta(S)} \int u_E\left(f\left(s\right)\right)dq\left(s\right)+\frac{1}{\lambda_E}R\left(q||p_E\right) \geq \min_{q\in\Delta(S)} \int u_E\left(f\right)dq+\frac{1}{\lambda_E}R\left(q||p_E\right)$$

where u is a nonconstant affine function, $\lambda_E \in [0,\infty)$, and $p_E \in \Delta^{\sigma}\left(S\right)$. Moreover u_E can be chosen to be the same for all such E and $\operatorname{supp} p_E \subseteq E$.

Proof The first part follows by Lemma 5 and Theorem 1 of (Strzalecki, 2011). For the second part, notice that by Theorem 3 of (Maccheroni, Marinacci, and Rustichini, 2006), $\succsim$ admits a variational representation:

$$f \succsim g \Longleftrightarrow \min_{q\in\Delta}\left(\int u\left(f\right)dq + c\left(q\right)\right) \geq \min_{q\in\Delta}\left(\int u\left(g\right)dq + c\left(q\right)\right)$$

for some nonconstant affine $u : X \rightarrow \mathbb{R}$ and a lower semicontinuous and grounded function $c : \Delta\left(S\right) \rightarrow [0,\infty]$. Notice that $\succsim$ and $\succsim_E$ coincide on X. Indeed, by definition of $\succsim_E$

$$x \succsim_E y \Rightarrow x \succsim yEx$$

and $yEx \succsim y$ by Weak Certainty Independence. Therefore, by Weak Order of $\succsim$, $x \succsim y$. Let $x \succsim y$, then by Theorem 3 of (Maccheroni, Marinacci, and Rustichini, 2006) $u\left(x\right) \geq u\left(y\right)$. Since c is grounded, there exists q^* with

$c(q^*) = 0$. But then

$$u(x) \geq u(y)\, q^*(E) + (1 - q^*(E))\, u(x) \geq \min_{q \in \Delta} \left(u(y)\, q(E) + (1 - q(E))\, u(x) + c(q) \right)$$

that is, $xEx \succsim yEx$, and $x \succsim_E y$. Therefore, by the Mixture Space Theorem (Herstein and Milnor, 1953), u and u_E are positive affine transformations of each other, concluding the proof. ■

We say that an act f is *unambiguously preferred* to g, denoted as $f \succsim^* g$ if

$$\alpha f + (1 - \alpha)\, h \succsim \alpha g + (1 - \alpha)\, h \text{ for all } \alpha \in [0, 1]\,, h \in \mathcal{F}.$$

Definition 6. A tuple (u, Q, μ, ϕ) is an average second order utility representation of the preference relation $\succsim$ if u is a nonconstant affine function, $\mu \in \Delta(\Delta(S))$, $Q \subseteq \Delta(S)$ is a nonempty set $\phi : u(X) \to \mathbb{R}$ is a strictly increasing continuous function and

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

Lemma 6. *(i) If $\succsim$ admits an average robust control representation then it admits an average second order utility representation. (ii) If $\succsim$ admits an average second order utility representation with $\phi(z) = -\exp(-z)$, then it admits an average robust control representation.*

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

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

Therefore,

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

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

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

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

Therefore,

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

■

Definition 7. A tuple $(u, \phi, \mathcal{G}, \pi)$ is a robust predictive representation of the preference relation $\succsim$ if u is a nonconstant affine function, $\phi : u(X) \to \mathbb{R}$ is a strictly increasing continuous function, $\mathcal{G}$ is a sub-sigma algebra of Σ, $\pi \in \Delta(S)$ is a probability measure nonatomic on $\mathcal{G}$ and

$$f \succsim g \Longleftrightarrow \int_S \phi^{-1}\left(\int_S \phi(u(f))\,\mathrm{d}\pi(s|\mathcal{G})\right)\mathrm{d}\pi \geq \int_S \phi^{-1}\left(\int_S \phi(u(g))\,\mathrm{d}\pi(s|\mathcal{G})\right)\mathrm{d}\pi.$$

A robust predictive representation requires that conditional to the information on a sigma algebra $\mathcal{G}$, the agents evaluate act f as subjective expected utility maximizer with function $\phi \circ u$.

For every $Q \subseteq \Delta(\Delta(S))$, and $\mu \in \Delta(Q)$, let $\mathcal{G}_Q = \{E \in \Sigma : q(E) \in \{0,1\} \text{ for all } q \in Q\}$, and let $\pi_\mu \in \Delta(S)$ be defined by

$$\pi_\mu(E) = \int_Q q(E)\,\mathrm{d}\mu(q).$$

Let $\phi(u) = -\exp(-\lambda u)$.

Lemma 7. *If* $\succsim$ *admits an identifiable average robust control representation* (u, Q, μ, λ), *then* $\succsim$ *admits the robust predictive representation* $(u, -\exp(-\lambda(\cdot)), \mathcal{G}_Q, \pi_\mu)$.

Proof Suppose that (u, Q, μ, λ) is an identifiable average robust control representation and that k identifies Q. By Lemma 25 in (Denti and Pomatto, 2020) π_μ is nonatomic on $\mathcal{G}_Q$. It is enough to show that for all $f \in \mathcal{F}$

$$\int_Q \min_{p \in \Delta(S)} \int_S \left[u(f) + \frac{1}{\lambda} R(p||q) \right] \mathrm{d}p \mathrm{d}\mu(q) = \phi^{-1} \left(\int_S \phi(u(f)) \mathrm{d}\pi_\mu(s|\mathcal{G}_Q) \right) \mathrm{d}\pi_\mu(q).$$

Let f be $\mathcal{G}_Q$-measurable and fix $q \in Q$. Then,

$$\{s : f(s) = \mathbb{E}_q(f)\}, \{s : f(s) > \mathbb{E}_q(f)\}, \{s : f(s) < \mathbb{E}_q(f)\} \in \mathcal{G}_Q,$$

and if $q(\{s : f(s) = \mathbb{E}_q(f)\}) = 0$, then either $q(\{s : f(s) > \mathbb{E}_q(f)\}) = 1$ or $q(\{s : f(s) > \mathbb{E}_q(f)\}) = 1$, a contradiction with the definition of $\mathbb{E}_q$. Hence, $\mathbb{E}_q\left[\int_S \phi^{-1}(f)\right] = \phi^{-1}(\mathbb{E}_q[f])$ which implies

$$\begin{aligned} \int_Q \phi^{-1} \left(\int_S f \mathrm{d}q \right) \mathrm{d}\mu(q) &= \int_Q \int_S \phi^{-1}(f) \mathrm{d}q \mathrm{d}\mu(q) \\ &= \int_S \phi^{-1}(f) \mathrm{d}\pi_\mu \\ &= \int_S \phi^{-1} \left(\int_S f \mathrm{d}\pi(s|\mathcal{G}_Q) \right) \mathrm{d}\pi_\mu. \end{aligned} \tag{2}$$

Then, let $f \in \mathcal{F}$. By Lemma 23 in (Denti and Pomatto, 2020)

$$\int_S \phi(u(f)) \mathrm{d}q = \int_S \int_S \phi(u(f)) \mathrm{d}k(s) \mathrm{d}q(s) \tag{3}$$

so that

$$
\begin{aligned}
\int_Q \min_{p\in\Delta(S)} \int_S \left[u(f)+\frac{1}{\lambda}R(p||q)\right] \mathrm{d}\mu(q) &= \int_Q \phi^{-1}\left(\int_S \phi(u(f))\,\mathrm{d}q(s)\right)\mathrm{d}\mu(q) \\
&= \int_Q \phi^{-1}\left(\int_S\int_S \phi(u(f))\,\mathrm{d}k(s)\,\mathrm{d}q(s)\right)\mathrm{d}\mu(q) \\
&= \int_S \phi^{-1}\left(\int_S \phi(u(f))\,\mathrm{d}k(s)\right)\mathrm{d}\pi_\mu \\
&= \int_S \phi^{-1}\left(\int_S \phi(u(f))\,\mathrm{d}\pi(s|\mathcal{G}_Q)\right)\mathrm{d}\pi_\mu
\end{aligned}
$$

where the the second inequality follows by equation (3), the third by the $\mathcal{G}_Q$-measurability of $s \mapsto \int_S \phi(u(f))\mathrm{d}k(s)$ and by equation (2), and the last equality follows by Lemma 24 in (Denti and Pomatto, 2020). ■

Lemma 8. *If* $\succsim$ *admits a robust predictive representation* $(u, -\exp(-\lambda(\cdot)), \mathcal{G}, \pi)$*, then* $\succsim$ *admits an identifiable average robust control representation* (u, Q, μ, λ) *where* $\pi = \pi_\mu$ *and* $\mathcal{G} = \sigma(k)$ *up to null events.*

Proof Suppose that $(u, \phi, \mathcal{G}, \pi)$ is a robust predictive representation of the preference relation $\succsim$. Let k be the regular conditional probability of π with respect to $\mathcal{G}$ whose existence is guaranteed by the fact that (S, Σ) is standard Borel. Define the prior $\mu \in \Delta(\Delta(\Omega))$ as the pushforward of π under k. Since k is a conditional probability with respect to $\mathcal{G}$, for each fixed $E \in \Sigma$ the functions $s \mapsto k(s, E)$ and $s \mapsto k(s, E)^2$ are $\mathcal{G}$-measurable and by definition of conditional probability

$$
\int_S k(s', E)\, k(s, \mathrm{d}s') = k(s, E) \ \text{ and } \int_S k(s', E)^2\, k(s, \mathrm{d}s') = k(s, E)^2 .
$$

So

$$
\begin{aligned}
\int_S k(s, E)^2\, \mathrm{d}p(s) + p(E)^2 &= \int_S\int_S k(s', E)^2\, k(s, \mathrm{d}s')\, \mathrm{d}p(s) + p(E)^2 \\
&= 2p(E)\int_S k(s, E)\, \mathrm{d}p(s)
\end{aligned}
$$

and therefore $p(\{s : k(s,E) = p(E)\}) = 1$ for μ almost all p. Let

$$Q = \{p : p(\{s : k(s) = p\}) = 1\}.$$

Since the space (S, Σ) is standard Borel, there exists a countable Σ' that generates Σ. We have just proved that for each $E \in \Sigma'$, for μ almost all p

$$p(\{\{s : k(s,E)\} s = p(E)\}) = 1$$

and so

$$p(\{s : k(s) = p\}) = 1.$$

Therefore, $\mu(Q) = 1$ and k is a $\mathcal{G} - \Sigma_\Delta^Q$ measurable function that identifies $\mathcal{P}$. With a change of variable, we get

$$\int_S \phi^{-1}\left(\int_S \phi(u(f)) \,\mathrm{d}\pi(\cdot|\mathcal{G})\right) \mathrm{d}\pi = \int_Q \phi^{-1}\left(\int_S \phi(u(f(s'))) \,\mathrm{d}p(s, \mathrm{d}s')\right) \mathrm{d}\mu(p).$$

But then by Lemma 6

$$\int_Q \phi^{-1}\left(\int_S \phi(u(f(s'))) \,\mathrm{d}p(s, \mathrm{d}s')\right) \mathrm{d}\mu(p) = \int_Q \min_{p \in \Delta(S)} \int_S \left[u(f) + \frac{1}{\lambda} R(p||q)\right] \mathrm{d}\mu(q).$$

Since the average multiplier representation that we have obtained is identifiable, it is enough to show that μ is nonatomic. Take any partition $E_1, ..., E_n$ of events that are equally likely under Σ, that exists since (S, Σ, π) is not atomic. For each $i \in \{1, .., n\}$ the set

$$E_i^* = \{p \in Q : q(E_i) = 1\}$$

has

$$\mu(E_i^*) = \mu(\{p \in Q : q(E_i) = 1\}) = \pi(\{s : k(s,E) = 1\}) = \pi(E_i) = \frac{1}{n}.$$

Proof of Lemma 1. We have

$$\begin{aligned}
\bar{G}_t^t(h_t, Q) &= \left(\log\left(\frac{\mathbb{P}_{p(h_t)}(h_t)}{\max_{\theta\in Q}\mathbb{P}_\theta(h_t)}\right)\right)/t \\
&= \left(\log\left(\frac{\prod_{i=1}^t p(h_t)(y_t)}{\prod_{i=1}^t q(y_t)}\right)\right)/t \\
&= \left(\log\left(\prod_{i=1}^t p(h_t)(y_t)\right) - \log\left(\prod_{i=1}^t q(y_t)\right)\right)/t \\
&= \left(\log\left(\prod_{y\in Y} p(h_t)(y)^{tp(h_t)(y)}\right) - \log\left(\prod_{y\in Y} q(y)^{tp(h_t)(y)}\right)\right)/t \\
&= \left(\sum_{y\in Y} tp(h_t)(y)\log(p(h_t)(y)) - \sum_{y\in Y} tp(h_t)(y)\log\left(\prod_{y\in Y} q(y)\right)\right)/t \\
&= R(p(h_t)||q).
\end{aligned}$$

Lemma 9. *Suppose that for any prior belief ν_0 supported on Θ and any optimal policy $\tilde{\pi}$ $\mathbb{P}_{\pi^b}[b = \tilde{\pi}(\nu_\tau)$ for all $\tau \geq 0] = 0$, then b is not a limit action.*

Proof Suppose by way of contradiction that there is an optimal policy $\tilde{\pi}$ and a history (a^t, y^t) with $\mathbb{P}_{\tilde{\pi}}[(a^t, y^t)] > 0$ such that with positive probability $\tilde{\pi}$ prescribes b after (a^t, y^t) in every future period. Define $\nu_0 = \mu(\cdot|(a^t, y^t))$, and notice that $\operatorname{supp}\nu_0 = \operatorname{supp}\mu_0 = \Theta$. Define ν_t to be the belief if the agent uses the policy π^b, i.e. plays b in every period. As the evolution of beliefs under π^b is the same as under $\tilde{\pi}$ for every history where the agent continues to play b, we have that $\mathbb{P}_{\tilde{\pi}}[b = \tilde{\pi}(\mu_\tau)$ for all $\tau \geq t] > 0$ if and only if $\mathbb{P}_{\pi^b}[b = \tilde{\pi}(\nu_\tau)$ for all $\tau \geq 0] > 0$. However, the later equals zero by the assumption of the lemma, which establishes that b cannot be a limit action.
■

Proof of Proposition 3. That each $\succsim^H$ admits an average robust control representation $(u_H, Q_H, \mu_H, \lambda_H)$ follows by Theorem 2. That $Q_H = Q$ follows from the representation and the fact that Q is the smallest closed set E

such $S \times (\Delta(S) \setminus E)$. That $u_h = u$ for some constant affine u follows from Constant Preference Invariance, and λ_H is increasing in LRT_Q by Proposition 8 in (Maccheroni, Marinacci, and Rustichini, 2006). ■

References

Cerreia-Vioglio, S., L. P. Hansen, F. Maccheroni, and M. Marinacci (2020). "Making decisions under model misspecification". *University of Chicago, Becker Friedman Institute for Economics Working Paper.*

Cerreia-Vioglio, S., F. Maccheroni, M. Marinacci, and L. Montrucchio (2013). "Classical subjective expected utility". *Proceedings of the National Academy of Sciences* 110, pp. 6754–6759.

Denti, T. and L. Pomatto (2020). *Model and Predictive Uncertainty: A Foundation for Smooth Ambiguity Preferences.* Tech. rep. Working paper.

Dupuis, P. and R. S. Ellis (2011). *A weak convergence approach to the theory of large deviations.* Vol. 902. John Wiley & Sons.

Esponda, I. and D. Pouzo (2016). "Berk–Nash equilibrium: A Framework for Modeling Agents with Misspecified Models". *Econometrica* 84, pp. 1093–1130.

— (2019). "Equilibrium in Misspecified Markov Decision Processes". *arXiv preprint arXiv:1502.06901.*

Folland, G. B. (1999). *Real analysis: modern techniques and their applications.* Vol. 40. John Wiley & Sons.

Frick, M., R. Iijima, and Y. Ishii (2020). "Stability and Robustness in Misspecified Learning Models".

Fudenberg, D., G. Lanzani, and P. Strack (2020). "Limits Points of Endogenous Misspecified Learning". *Available at SSRN.*

Gilboa, I., D. Schmeidler, et al. (1989). "Maxmin expected utility with non-unique prior". *Journal of Mathematical Economics* 18, pp. 141–153.

Hansen, L. P. and T. J. Sargent (2020). *Structured Ambiguity and Model Misspecification*.

— (2001). "Robust control and model uncertainty". *American Economic Review* 91, pp. 60–66.

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

Maccheroni, F., M. Marinacci, and A. Rustichini (2006). "Ambiguity aversion, robustness, and the variational representation of preferences". *Econometrica* 74, pp. 1447–1498.

Savage, L. J. (1972). *The foundations of statistics*. Courier Corporation.

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.