

# Robust Comparative Statics with Misspecified Bayesian Learning\*

Aniruddha Ghosh<sup>†</sup>

This Version: December 8, 2024

## Abstract

We present novel monotone comparative statics results for steady state behavior in a dynamic optimization environment with misspecified Bayesian learning. We consider a generalized framework, based on Esponda and Pouzo (2021), wherein a Bayesian learner facing a dynamic optimization problem has a prior on a set of parameterized transition probability functions (models) but is misspecified in the sense that the true process is not within this set. In the steady state, the learner infers the model that best-fits the data generated by their actions, and in turn, their actions are optimally chosen given their inferred model. We characterize conditions on the primitives of the environment, and in particular, over the set of models under which the steady state distribution over states and actions and inferred models exhibit monotonic behavior. Further, we offer a new theorem on the existence of a steady state on the basis of a monotonicity argument. Lastly, we provide an upper bound on the cost of misspecification, again in terms of the primitives of the environment. We demonstrate the utility of our results for several environments of general interest, including forecasting models, dynamic effort-task, and optimal consumption-savings problems.

*Journal of Economic Literature* Classification Numbers: C1, C5, C6, D8, D9

*Key Words:* Bayesian learning, Berk-Nash equilibrium, comparative statics, misspecification.

---

\*For numerous comments and guidance during the course of the project, I would like to thank my committee members Robert Anderson, Robert Barbera, Ying Chen, and M. Ali Khan. I am also thankful to Chris Carroll, Yujian Chen, Liuchun Deng, Haosui Duanmu, Marcelo Fernandez, Mira Frick, Ashwin Kambhampati, Edi Karni, Jong Jae Lee, Andrew Mackenzie, Vikram Manjunath, Arina Nikandrova, Mallesh Pai, Arthur Paul Pedersen, Metin Uyanik, Marie-Louise Vierø, Xinyang Wang, and Sezer Yasar for their comments and encouragement at various stages. A version of the results reported here were first presented at SAET 2022, and subsequent elaborations were presented at SUNY 2023 and the departmental ‘economic theory seminar’ at Johns Hopkins University. All remaining errors are my own.

<sup>†</sup>Department of Economics, Cal Poly, San Luis Obispo, CA 93401. **E-mail** aghosh10@calpoly.edu.

## Contents

1	Introduction	2
2	General Framework	8
3	Motivating Examples	15
4	Main Results	19
5	Analysis of Examples	24
6	Welfare Comparisons	26
7	Concluding Remarks and Extensions	28
A	Mathematical Preliminaries	30
B	Proofs	31
	References	38

*In general, dynamic programs with learning are intractable, i.e., they are not solvable either analytically or numerically, when there is no separability of control and learning. The problem is not only whether a solution exists, but if a solution can be characterized and its properties studied.*

Koulovitanos, Mirman, and Santugini (2009)

## 1 Introduction

Stemming from the foundational contributions of Arrow and Green (1973) and Nyarko (1991), model misspecification in economic environments continues to be of substantive interest for economists. It acknowledges the possibility that economic agents, perhaps due to cognitive limitations or simplified perspectives, may often not know the true model of their complex economic environment and, therefore, learn about it through a set of models that are *misspecified* in the sense that the true model is not contained in their set of models. A growing body of literature in economics suggests that enriching classical economic environments with misspecified economic agents aligns theoretical predictions closely with the observed behavior.<sup>1</sup>

Dynamic programming offers one robust framework for economists to approach discrete-time economic problems. They are commonly used to analyze instances where economic agents make decisions sequentially in an uncertain environment. The framework applies to a broad spectrum of economic models, from consumption-savings decisions for individuals to investment choices for firms. However, they are known to be intractable, no closed-form solutions, when incorporated with agents who simultaneously make decisions and learn about their environment through their models, where both the decision making and the learning activity mutually affect each other. In this paper, we provide *robust* predictions on the comparative statics properties of an agent's behavior within a dynamic programming environment, particularly when they are learning with potentially misspecified models.<sup>2</sup>

Consider the following setting, conceptualized by Esponda and Pouzo (2021) (hereafter, EP), in the context of a single-agent dynamic optimization problem, a Markov Decision Process (hereafter, MDP). In each period  $t = 0, 1, 2, \dots$ , an agent observes a state  $s_t$  and then takes an action  $x_t$  and receives a utility,  $u(s_t, x_t)$ . The current action and the state together determine the evolution of the state tomorrow,  $s_{t+1}$ , via the true transition function,  $Q(\cdot|s_t, x_t)$ . The agent chooses a sequence of actions to maximize their current and the expected discounted flow of utility. The agent doesn't know the true (objective) transition function and, therefore, chooses actions based on a set of model (subjective) transition functions,  $\{Q_\theta, \theta \in \Theta\}$ , parameterized by

---

<sup>1</sup>For example, Farmer, Nakamura, and Steinsson (2023) finds that much of the anomalies concerned with forecasts of professional forecasters can be attributed to them not knowing the true model of the environment.

<sup>2</sup>As is well-known, comparative statics analysis deals with characterizing conditions for various environments, both static and dynamic, under which decision rules and solution concepts are increasing in the primitives.

$\theta$ , that is misspecified in the sense that it does not contain the true transition function. The Bayesian agent has a prior  $\mu$  on the set of models  $\Theta$  and, every period updates the prior using Bayes' rule  $B$ , based on the current state, the chosen action, and the realized state. The agent's problem can be formulated as a Bellman equation,

$$V(s, \mu) = \max_{x \in \mathbb{X}} \left\{ u(s, x) + \beta \int_{\mathbb{S}} V(s', \mu') \bar{Q}_\mu(ds' | s, x) \right\}, \quad (1)$$

where  $\bar{Q}_\mu = \int_{\Theta} Q_\theta \mu(d\theta)$ ,  $\mu' = B(s, x, s', \mu)$  is the next period's belief,<sup>3</sup> updated using Bayes' rule, and  $V : \mathbb{S} \times \Delta(\Theta) \rightarrow \mathbb{R}$  is the unique solution to the Bellman equation of the agent, where  $\Delta(\Theta)$  is the set of probability distributions on the parameter space,  $\Theta$ .

For a misspecified agent that follows the above setting in choosing their sequence of actions and updating their sequence of beliefs over models, EP predicts its steady-state behavior in terms of a solution concept called the *Berk-Nash equilibrium*. It is an equilibrium distribution over states and actions corresponding to which the misspecified agent infers the model (or a set of models) that 'best fits' the true model. In turn, given the inferred model, the chosen actions are optimal for each state. Furthermore, the equilibrium distribution over states and actions gives rise to a stationary Markov process. The 'best-fit' formalization is in the sense of a minimum weighted Kullback-Liebler divergence, with weights determined by the equilibrium distribution over states and actions.<sup>4</sup> Given this prediction of the long-run behavior of a misspecified Bayesian agent in terms of a Berk-Nash equilibrium and its corresponding inferred model, we ask the following comparative statics question: how do the equilibrium objects respond to changes in the economic primitives? Phrased differently, what are the requirements on the primitives of the economic environment so that a misspecified agent's stationary distribution and corresponding best-fit models exhibit monotone comparative statics properties?

This question is of general theoretical interest for several reasons. First, models of dynamic programming incorporated with learning, whether correctly specified or misspecified, such as the one described in Equation (1), are known to be intractable.<sup>5</sup> In an important paper on learning in stochastic growth models, Koulovatianos, Mirman, and Santugini (2009) (hereafter, KMS) note that the intractability arises primarily for two reasons: (a) the curse of dimensionality problem, which occurs because of incorporation of beliefs about the models in the state space,

---

<sup>3</sup>The Bayesian operator  $B$ , where for any  $A \subseteq \Theta$ ,  $B(s, s', x, \mu)(A) = \int_A Q_\theta(s'|s, x) \mu(d\theta) / \int_{\Theta} Q_\theta(s'|s, x) \mu(d\theta)$ .

<sup>4</sup>The Kullback-Liebler (KL) divergence measures the difference between two probability distributions in terms of relative entropy; in our case, between the models and the true process; see Cover and Thomas (2005) for details.

<sup>5</sup>With learning incorporated in dynamic programming models, the agent is an active learner and processes incoming data in order to infer the models of their unknown environment. At the same time, they also take decisions. These two functions of the agent are intertwined.

thus making the state space very large and computationally cumbersome, and (b) with Bayesian updating, the prior and the posterior over models potentially belong to different families, thereby hampering analytical and computational tractability.<sup>6</sup> By focusing on the long-run prediction for a misspecified dynamic program with a precise characterization of the distribution over states and actions, as well as the limiting posterior (beliefs) over models, we abstract away from the intractability problems and thereby make predictions about their steady-state properties even when their closed-form solutions may not exist. Thus, in this paper, we expand the theory of dynamic programming with misspecified learning, potentially rendering it applicable across a wide spectrum of economic domains.

Second, comparative statics for misspecified learning environments are typically not well-understood in the existing literature on misspecification in economic theory. This traces itself to the issue that with concurrent decision making and learning, there is often no separability between the two, as is illustrated in Equation (1). We tackle this problem by recognizing that steady states for such models are fixed points of a well-defined equilibrium mapping. This insight allows us to appeal to a relatively newer literature of powerful order-theoretic techniques, tailored for dealing with the monotonicity of fixed points for non-lattice spaces (Acemoglu and Jensen (2015)). A technical challenge that lies with misspecified MDPs is that the underlying spaces containing such fixed points are typically not lattices in any natural order, and therefore, a wide array of popular lattice-dependent techniques (Hopenhayn and Prescott (1992); Topkis (1998)) are not applicable. Finally, our results for comparative statics do not necessitate any specific knowledge of the functional forms of any environment primitives, and therefore are *robust*, and hold for a very general class of dynamic programming environments with misspecified models.<sup>7</sup>

To interpret our framing, we discuss examples that cover three important economic environments. The first example concerns Bayesian inference with misspecified AR(1) models. The agent’s set of models is misspecified in the sense that although the true process is AR(1), the innovations of the process are non-normally distributed and, therefore, are not contained in their set of AR(1) models with Gaussian innovations. We characterize the Berk-Nash equilibrium for this pathological example, and show that despite being misspecified, the Bayesian agent correctly infers the persistence parameter of the true process at the steady state. The next two examples are based on EP, albeit with slight modifications for interpretation and analytical tractability. In the second example, we study a canonical problem of effort provision in a dynamic effort-task

---

<sup>6</sup>As KMS and EP note, the curse of dimensionality does not arise if the beliefs are over a finite parameter set. However, for parameter spaces such as the real line, it is a cause for concern. Also note that, unlike KMS, in our framework, choosing actions also affects the flow of information, thus further complicating matters.

<sup>7</sup>The reader may be interested in the following discussion: <https://stats.stackexchange.com/questions/274815/why-should-i-be-bayesian-when-my-model-is-wrong>. Uppal and Wang (2003) discuss the contrast between Savagian and Knightian approaches to tackle model misspecification concerns. The interested reader is referred to their p. 2469 for further reading.

problem where the agent seeks to learn their ability while exercising control over the stochastic outcome of a task by choosing their effort levels. The agent’s models are misspecified since they incorrectly postulate that tomorrow’s outcome depends only on their ability and not on today’s outcome. In contrary, the true dynamics are determined by both ability and current outcome. The Berk-Nash equilibrium for this case is the steady-state frequency of the outcome, either a success or a failure, along with the agent’s inference of their ability. We characterize the dependence of these equilibrium objects on the economic primitives of the environment, such as the cost of exerting higher effort or the objective probability of success, and study their associated comparative statics. The final example considers an agent that optimally chooses its consumption and savings in the presence of preference and productivity shocks. The agent seeks to learn about the return of their wealth process but does so through a set of misspecified models that do not consider the correlation between the two shocks. Here, the Berk-Nash equilibrium is the long-run perceived distribution of the wealth process along with their corresponding inferred parameter of the return on that process. However, this case has no analytical solution; therefore, making our results most valuable for such settings.

We now turn to the specificity of our results. Theorem 1 establishes the existence of a Berk-Nash equilibrium for a misspecified MDP with compact action and state spaces. EP and Anderson, Duanmu, Ghosh, and Khan (2023) (hereafter, ADGK) provide proofs for the existence of a Berk-Nash equilibrium with finite and infinite state and actions, respectively. ADGK applies recent advances in non-standard analysis to present existence theorems for misspecified environments with infinite action and state spaces and unbounded payoff functions. We adopt the more general environment of ADGK, but the novelty of our existence result lies in exploiting the non-lattice structure of our economic environment. In general, distributions over states and actions are not lattices under most natural orders; therefore, to show the existence of the Berk-Nash equilibrium which in turn is a distribution, we rely on a novel application of techniques tailored for non-lattice spaces. These techniques, originating in Smithson (1971), have been pioneered in Acemoglu and Jensen (2015) for establishing the existence and comparative statics results in infinite-horizon large dynamic economies.

After furnishing the existence result for the Berk-Nash equilibrium, we utilize these techniques for Theorems 2-4 and give sufficient conditions for which the Berk-Nash equilibrium and the corresponding inferred best-fit model respond monotonically to the primitives of the environment. We formalize a notion of a positive shock for a misspecified dynamic optimization problem, defined as shocks to the primitives that increase strategies for any given beliefs over models, and under mild assumptions establish in Theorem 2 that the least and the greatest inferred best-fit model at the steady state, responds monotonically to a change in the economic primitives. Theorem 2 further provides a prediction for the monotonicity of the Berk-Nash equi-

librium in the usual stochastic order, when a misspecified agent is hit with positive shocks to their primitives. In Theorem 3, we show that the least and the greatest inferred best-fit model responds monotonically to an *increase* in the set of models, where an increase is defined in the strong-set order in the space of parameters. Theorem 4 drops the usual stochastic order on the underlying space and extends Theorem 2 to spaces with increasing and convex orders.<sup>8</sup>

Our final result is of a slightly different flavor from the previous results. In Theorem 5, we perform a welfare comparison of steady-state learning in a correctly specified vis-à-vis a misspecified learning environment. While welfare under misspecified learning is always weakly lower than that under correct learning, we provide an upper bound on the cost of misspecification, formalized as the difference between the two instances and outline its dependence in terms of the primitives of the environment. We also note the potential usefulness of this bound for computational applications concerning Berk-Nash equilibria.

**Related literature.** This paper belongs to the growing literature on learning with misspecified models and its implications on the properties of the learning behavior.

Within the dynamic programming framework, following EP and ADGK, this paper adds to this literature a more general framework to model MDPs with misspecified learning for infinite environments.<sup>9</sup> MDPs are commonly used to analyze instances involving agents making decisions sequentially in an uncertain environment and, therefore, have found many applications within natural and social sciences. Within the economics literature, one of its earliest uses is traceable in Arrow, Harris, and Marschak (1951), where an inventory holder maximizes some given objective (profits, revenue, net utility) by choosing an inventory policy subject to stochastic product demand and other random fluctuations to the primitives of the problem.<sup>10</sup> Since then, MDPs have spanned both micro and macro environments, covering many important economic applications; these include (i) investment with adjustment costs under uncertain demand (Lucas and Prescott (1971)), (ii) one-sector optimal economic growth with uncertainty (Brock and Mirman (1972)), (iii) equilibrium search and unemployment (Lucas and Prescott (1974)), and (iv) asset prices in a pure exchange economy (Lucas (1978)).<sup>11</sup> In a recent paper, Saghaian (2018) presents a framework for an MDP in which agents possess a ‘cloud’ of models and allows for considering model ambiguity and the potential misspecification within an MDP.

---

<sup>8</sup>As Che, Kim, and Kojima (2021) note in the context of individual choices, the strong set order proves to be an appropriate notion; see the references therein for further details.

<sup>9</sup>See Puterman (1994) for illustrations covering operations research, engineering, and many other allied fields. Also, see Rust (1994) for an extensive survey on MDPs and associated structural estimation methods. The results in this paper could potentially be admissible in an extended learning version of the partially observable MDP framework of Saghaian (2018); see the references therein for an overview of that literature.

<sup>10</sup>This work has important precursors in Arrow, Blackwell, and Girshick (1949).

<sup>11</sup>Other prominent examples include models of portfolio choice under uncertainty (Phelps (1962); Levhari and Srinivasan (1969)) and business cycle models (Kydland and Prescott (1982); Long Jr. and Plosser (1983)).

This paper also contributes to the understanding of the comparative static properties of the limit posterior of a Bayesian inference process. In his foundational paper on inference with misspecified models, devoid of any actions, Berk (1966) demonstrates that for a misspecified Bayesian agent learning about a parameter through a series of independently and identically distributed (i.i.d.) signals, the posterior concentrates asymptotically on those set of models (referred therein as the asymptotic carrier) where the Kullback-Liebler divergence is minimal with respect to the true model.<sup>12</sup> In this paper, we address a misspecified MDP, a non-i.i.d. environment, that combines decision-making and inference but heavily relies on Berk’s characterization to abstract away from the dynamics of the updating process. This approach enables us to focus on the comparative statics of the steady-state behavior for both the Berk-Nash equilibrium and the corresponding asymptotic carrier, and contribute to the inference literature by outlining the dependence of the inferred models on the primitives of the decision making and the learning environment. Furthermore, our results also have implications for a series of papers that follow the papers of Huber (1967) and White (1982). These papers employ maximum likelihood techniques for inference and show that when an agent conducts inference using maximum likelihood estimation, their inference converges towards models that minimize the Kullback–Leibler divergence with respect to the true distribution, identical to the limits proposed in Berk (1966).<sup>13</sup>

The topic of misspecification in economic environments has been a subject of active research area.<sup>14</sup> In an influential paper, Esponda and Pouzo (2016) provides a framework for a static game-theoretic setting that relaxes the assumption that agents have a correct view of the game’s environment. This formulation is further extended to the dynamic programming environment in EP for a finite (states and actions) setting and by ADGK for an infinite setting. While there has been influential work on social learning problems as in Frick, Iijima, and Ishii (2020) and Bohren and Hauser (2021), properties of asymptotic learning in a misspecified MDP environment are still not well-studied. The results in this paper contribute to the understanding of such properties, in terms of the comparative statics behavior of misspecified agents.

Naturally, this work is informed by the substantive literature on monotone comparative statics. The theory of monotone comparative statics deals with characterizing conditions under which the optimizing behavior of agents leads to solution concepts (and invariant distributions) being monotonic in the primitives of the environment.<sup>15</sup> In an influential paper, Hopenhayn and

---

<sup>12</sup>See also Bunke and Milhaud (1998) and Shalizi (2009) for extensions to non-iid environments. The reader should note that Berk’s asymptotic carrier is independent of the prior distribution.

<sup>13</sup>For Bayesian inference, Grünwald and van Ommen (2017) classifies misspecification into good or bad, depending on the structure of the models. They further show that this has important implications for the consistency of Bayesian inference.

<sup>14</sup>Among many other prominent works, some notable ones include Kirman (1975), Jehiel (2005), and Hansen and Sargent (2011). The reader is referred to the references therein for a broader view of the literature.

<sup>15</sup>Some important references include Milgrom and Shannon (1994), Amir (1996), Topkis (1998), Huggett (2003),



Prescott (1992) uses the tools of monotone comparative statics to analyze stationary dynamic optimization problems with lattice environments. These results turn out to be of limited use for obtaining our theorems. Our proofs instead exploit the non-lattice structure of our setting and, therefore, depend on the techniques developed in Smithson (1971) and Acemoglu and Jensen (2015). We adapt and apply these techniques in a novel way to misspecified MDPs with Bayesian learning.

Lastly, we calculate the cost of misspecification in terms of the discrepancy between the expected discounted welfare under correctly specified and misspecified Bayesian learning for an MDP and provide an upper bound on the discrepancy between the two quantities. We then characterize its comparative statics properties with respect to the primitives. This is partly inspired by Santos (2000)’s work on testing the accuracy of numerical solutions in dynamic models. We argue that the one-way bound on welfare that we provide is useful for computational purposes of such equilibria.

**Outline.** The paper is organized as follows. Section 2 sets up the general framework for a misspecified MDP and outlines the necessary prerequisites and techniques for the comparative statics analysis. Section 3 offers three examples to illustrate our setting. The main results of this paper are presented in Section 4, while Section 5 completes the formal analysis of the examples. Section 6 compares welfare between correctly specified and misspecified settings. Section 7 concludes by discussing how these results can be extended and related to different settings. Appendix A to this paper contains the necessary mathematical preliminaries, and the proofs of the results are presented in Appendix B.

## 2 General Framework

We begin by framing the environment for a Markov Decision Process (MDP) with model misspecification. While the conceptual framework is patterned after the finite (states and actions) environment of EP, we adopt the setting of ADGK since it applies to infinite settings; an environment that naturally features many important applications in economics. After outlining the contents of a misspecified MDP and the relevant Berk-Nash equilibrium concept, we provide an overview of the order-theoretic methods that are instrumental for our results. A refresher for these concepts and methods is provided in Appendix A for the reader’s convenience.

At the start of each period  $t = 0, 1, 2, \dots$ , the agent observes a state realization,  $s_t \in \mathbb{S}$ , and then chooses an action,  $x_t \in \mathbb{X}$ . Given a transition probability function  $Q(\cdot | s_t, x_t)$ , the state and action together determine the distribution of the next period state,  $s_{t+1}$ . The per-period

---

Datta, Reffett, and Woźny (2018), and Light (2021). Amir (2018) is an useful guide for some of the history and also of the more recent advances in the comparative statics literature.

payoff function is a mapping  $u : \mathbb{S} \times \mathbb{X} \rightarrow \mathbb{R}$ . The agent maximizes expected discounted utility (discount factor  $0 < \beta < 1$ ) by choosing a feasible sequence of policy functions  $\{x_t\}_{t=1}^{\infty}$  that solves the following problem,

$$V(s_0) = \max_{\{x_t\}_{t=0}^{\infty}} \mathbb{E}_Q \left[ \sum_{t=0}^{\infty} \beta^t u(s_t, x_t) \right], \quad t = 0, 1, 2, \dots \quad (2)$$

Following (2), the Bellman equation for the agents' sequential decision problem is formulated as,

$$V(s) = \max_{x \in \mathbb{X}} \left\{ u(s, x) + \beta \int_{\mathbb{S}} V(s') Q(ds' | s, x) \right\}, \quad (3)$$

where  $V$ , the value function, is the unique solution to (3). Corresponding to this  $V$ , the optimal policy correspondence  $G$  is given by,

$$G(s) \equiv \operatorname{argmax}_{x \in \mathbb{X}} \left\{ u(s, x) + \beta \int_{\mathbb{S}} V(s') Q(ds' | s, x) \right\}, \quad (4)$$

where action  $\hat{x}$  is optimal given state  $s$ . We summarize this environment in the following definition of a MDP.<sup>16</sup>

**Definition 1.** A Markov Decision Process (hereafter, MDP) is a tuple  $\langle \mathbb{S}, \mathbb{X}, q_0, Q, u \rangle$ , where (i) the state space  $\mathbb{S}$  is a compact metric space with Borel  $\sigma$ -algebra  $\mathcal{B}(\mathbb{S})$ , (ii) the action space  $\mathbb{X}$  is a compact metric space with Borel  $\sigma$ -algebra  $\mathcal{B}(\mathbb{X})$ , (iii) the initial distribution of states  $q_0$  is a probability measure on  $\mathbb{S}$ , (iv)  $Q : \mathbb{S} \times \mathbb{X} \rightarrow \mathcal{M}_1(\mathbb{S})$  is a transition probability function, where  $\mathcal{M}_1(\mathbb{S})$  denote the set of probability measures on  $\mathbb{S}$ , and (v)  $u : \mathbb{S} \times \mathbb{X} \rightarrow \mathbb{R}$  is the per-period payoff function.<sup>17</sup>

We next define a Subjective Markov Decision Process. It adds to the tuple, a set of subjective transition functions,  $\{Q_{\theta}\}_{\theta \in \Theta}$ , parameterized with  $\theta$ . We refer to the parameter space  $\Theta$  as the set of *models*.

**Definition 2.** A subjective Markov Decision Process (hereafter, SMDP) is an MDP,  $\langle \mathbb{S}, \mathbb{X}, q_0, Q, u \rangle$ , and a non-empty family of transition probability functions,  $\mathcal{Q}_{\Theta} = \{Q_{\theta} : \theta \in \Theta\}$ , where each transition probability function  $Q_{\theta} : \mathbb{S} \times \mathbb{X} \rightarrow \mathcal{M}_1(\mathbb{S})$  is indexed by a parameter value  $\theta \in \Theta \subseteq \mathbb{R}$ . A SMDP is said to be *misspecified* if  $Q \notin \mathcal{Q}_{\Theta}$ .

Notice that under this definition, an SMDP could be misspecified in several ways. For instance, the true transition function and the set of model transition functions may pertain to dissimilar

<sup>16</sup>The framework can be extended with minor modifications to allow for the dependence of feasible set of actions on the state variable; see the Esponda and Pouzo (2015) working paper for details.

<sup>17</sup> $\mathcal{M}_1(\mathbb{S})$  is the space of finite probability measures on  $(\mathbb{S}, \mathcal{B}(\mathbb{S}))$  endowed with the weak-\* topology.

families of probability distributions. Furthermore, even if the true model and set of models belong to the same family of distributions, they can be misspecified if the support, the range of possible values of the parameter, of the models is different from that of the true distribution.<sup>18</sup> Our next definition requires the primitives of the SMDP to satisfy certain *regularity* conditions.

**Definition 3.** A regular SMDP satisfies the following conditions.

- (i) (Continuity) The mappings  $(s, x) \rightarrow Q(\cdot | s, x)$  and  $(\theta, s, x) \rightarrow Q_\theta(\cdot | s, x)$  are continuous in the Prokhorov metric, and the density function  $D_\theta(s' | s, x)$  is jointly continuous on the set

$$\{(\theta, s', s, x) : Q(s, x) \text{ is dominated by } Q_\theta(s, x)\},$$

where  $(D_\theta(s' | s, x))$  is the Radon-Nikodym derivative of  $Q$  with respect to  $Q_\theta$ ,  $\theta \in \Theta$ .

- (ii) (Absolute continuity) There is a dense set  $\hat{\Theta} \subset \Theta$  such that  $Q(\cdot | s, x)$  is absolutely continuous with respect to  $Q_\theta(\cdot | s, x)$  for all  $\theta \in \hat{\Theta}$  and  $(s, x) \in \mathbb{S} \times \mathbb{X}$ .
- (iii) (Uniform integrability) For every compact set  $S' \subset S$ , there exists some  $r > 0$  such that  $(D_\theta(\cdot | s, x))^{1+r}$  is uniformly integrable with respect to  $Q_\theta(\cdot | s, x)$  over the set  $\hat{\Theta}$ .<sup>19</sup>
- (iv) (Compactness) The parameter space  $\Theta$  is a compact metric space.

Condition (i) is a standard technical condition and requires the transition functions, both true and subjective, to be continuous. Condition (ii) requires that there always exists some model  $Q_\theta$  that accounts for every observation from the true distribution,  $Q$ . Condition (iii) places an uniform integrability requirement to deal with distributions with infinite support. Lastly, condition (iv) requires the parameter space  $\Theta$  to be compact. This is an essential requirement for the existence of best-fit models as will be evident in the next definition. We next measure the cost of misspecification in terms of the well-known measure of relative entropy, the Kullback-Liebler divergence.

**Definition 4.** For a given true transition function  $Q$  and model transition function  $Q_\theta$ , the Kullback-Liebler divergence,  $\mathcal{D}_{\text{KL}}$ , of  $Q_\theta$  with respect to  $Q$  is defined as,

$$\mathcal{D}_{\text{KL}}(Q(s, x), Q_\theta(s, x)) = \mathbb{E}_{Q(\cdot | s, x)} [\ln (D_\theta(s' | s, x))]. \quad (5)$$

---

<sup>18</sup>See Examples 1 and 2 for an instance for the former, and Example 3 for the latter. Den Haan and Drechsel (2019) explore more on forms of model misspecification and their relevance to econometric methods with macroeconomic models.

<sup>19</sup>The density function  $D_\theta(\cdot | s, x)$  is the Radon-Nikodym derivative of  $Q$  with respect to a model,  $Q_\theta$ . The uniform integrability is satisfied if the density functions  $D_\theta(\cdot | s, x)$  are uniformly bounded over the set  $\{(\theta, s, x) : Q(\cdot | s, x) \text{ is dominated by } Q_\theta(\cdot | s, x)\}$ . For example, given any two Gaussian distributions with distinct variances, the Radon-Nikodym derivative of the one with the larger variance with respect to the other is unbounded.

Then for any distribution over states and actions,  $m \in \mathcal{M}_1(\mathbb{S} \times \mathbb{X})$ , and model parameter,  $\theta \in \Theta$ , the *weighted* Kullback-Leibler divergence is a mapping  $K_Q : \mathcal{M}_1(\mathbb{S} \times \mathbb{X}) \times \Theta \rightarrow \bar{\mathbb{R}}_{\geq 0}$  such that

$$K_Q(m, \theta) = \int_{\mathbb{S} \times \mathbb{X}} \mathcal{D}_{\text{KL}}(Q(s, x), Q_\theta(s, x)) m(ds, dx).^{20} \quad (6)$$

The set of closest parameter values given a distribution  $m \in \mathcal{M}_1(\mathbb{S} \times \mathbb{X})$  and for a given true transition function  $Q$  is the set,  $\Theta(m; Q) = \underset{\theta \in \Theta}{\operatorname{argmin}} K_Q(m, \theta)$ .<sup>21</sup>

That is, given a distribution  $m$  over states and actions and a true transition function  $Q$ , the *best-fit* set  $\Theta(m; Q)$  is the set of those parameter values that minimize the weighted relative entropy between true transition function and the parameterized model transition functions in  $\Theta$ . Our focus in this paper is on the Berk-Nash equilibrium. We now define it.

**Definition 5.** A probability distribution  $m^* \in \mathcal{M}_1(\mathbb{S} \times \mathbb{X})$  is a Berk-Nash equilibrium of the regular-SMDP if there exists a belief  $\mu^* \in \mathcal{M}_1(\Theta)$  such that the following conditions hold.

- (a) For all states and actions,  $(s, x)$ , that are in the support of  $m^*$ , action  $x$  is *optimal* given state  $s$  in the MDP( $\bar{Q}_{\mu^*}$ ), where  $\bar{Q}_{\mu^*} = \int_{\Theta} Q_\theta \mu^*(d\theta)$ .
- (b) For a given true transition function  $Q$ , beliefs  $\mu^*$  are *restricted* over the best-fit set of parameters, that is,  $\mu^* \in \mathcal{M}_1(\Theta(m^*; Q))$ .
- (c) For all  $A \in \mathcal{B}(\mathbb{S})$ ,  $m_{\mathbb{S}}^*(A) = \int_{\mathbb{S} \times \mathbb{X}} Q(A|s, x) m^*(ds, dx)$ , where  $m_{\mathbb{S}}^*$  denote the *invariant* marginal measure of  $m^*$  on  $\mathbb{S}$ .

The Berk-Nash equilibrium is a probability distribution  $m^*$  over the states and actions, supported by equilibrium beliefs  $\mu^*$  over the best-fit set of parameterized models. The beliefs  $\mu^*$  are optimal given the data generated by the steady-state distribution over states and actions,  $m^*$ . In turn, the distribution  $m^*$  over states and actions is such that actions, conditioned on the state, are subjectively optimal given the equilibrium beliefs  $\mu^*$  over the best-fit models. For instance, in the case of a correctly specified MDP, where the true transition function  $Q$  is part of the set of models  $\mathcal{Q}_\Theta$ , the concept of Berk-Nash equilibrium implies that in the steady state for a Bayesian learner, beliefs concentrate on the true transition function, and the actions are optimal given this belief. The equilibrium concept effectively reduces to the well-known stationary solution for a MDP( $Q$ ).

<sup>20</sup> $\bar{\mathbb{R}}$  denotes the extended real line, equipped with the one-point compactification topology.

<sup>21</sup> $\mathcal{M}_1(\mathbb{S} \times \mathbb{X})$  is the space of finite probability measures on  $(\mathbb{S} \times \mathbb{X}, \mathcal{B}(\mathbb{S} \times \mathbb{X}))$  endowed with the weak-\* topology. The density function  $D_\theta(s'|s, x)$  is jointly continuous on the set  $\{(\theta, s', s, x) : Q(s, x) \text{ is dominated by } Q_\theta(s, x)\}$ . Further,  $\frac{0}{0} = 0$ ,  $\frac{1}{0} = \infty$ ,  $\log \infty = \infty$ . We follow the standard convention in that  $\ln(0) \cdot 0 = 0$  and integral of infinity over a set of measure 0 is 0.

However, if the SMDP is misspecified, that is, if  $Q$  is not contained within the set  $\mathcal{Q}_\Theta$ , then in Definition (5), condition (a) asserts that action  $x$  is optimal in the MDP( $\bar{Q}_{\mu^*}$ ) when the equilibrium transition function is a weighted combination of model transition functions, with weights given by  $\mu^*$ . Condition (b) mandates that these beliefs  $\mu^*$  are determined by minimizing the weighted relative entropy between the true transition function and the model transition functions, where the weights are the Berk-Nash equilibrium,  $m^*$ . Lastly, condition (c) stipulates that the marginal distributions over states,  $m_s^*$  is an invariant distribution. Notice that in condition (c), the invariant distribution is dependent on the optimal actions  $x$  via the true transition function,  $Q$ . EP show that under moderate conditions, the learning illustration with misspecified models sketched in Equation (1) converges to the Berk-Nash equilibrium.<sup>22</sup> The equilibrium serves as a prediction of the steady-state behavior in where  $m^*$  is the steady-state distribution over states and actions, and  $\mu^*$  is the limit of the sequence of posteriors of a Bayesian learner with misspecified models.

Our focus is on the comparative statics of the Berk-Nash equilibrium and the associated best-fit set with respect to the primitives  $P$  of the environment, where  $P = \langle u, \beta, Q, \mathcal{Q}_\Theta, \Theta \rangle$  are our objects in the collection of primitives. Given any primitive  $p \in P$ , we next adjust our value function, the optimal policy correspondence, and the best-fit set to allow for their dependence on the primitive. At the steady state solution, agent has a belief  $\mu^*$  over their set of best-fit models and solves a stationary dynamic programming problem where the value function  $V : \mathbb{S} \times \mathcal{M}_1(\Theta) \times P \rightarrow \mathbb{R}$  is determined by the following functional equation, where  $\bar{Q}_{\mu^*} = \int_{\Theta} Q_\theta \mu^*(d\theta)$ ,

$$V(s, \mu^*, p) = \max_{x \in \mathbb{X}} \left\{ u(s, x) + \beta \int_{\mathbb{S}} V(s', \mu^*, p) \bar{Q}_{\mu^*}(ds' | s, x) \right\} \text{ and} \quad (7)$$

corresponding to  $V$  and  $\mu^*$ , the stationary optimal policy correspondence  $G$  is given by,

$$G(s, \mu^*, p) \equiv \operatorname{argmax}_{x \in \mathbb{X}} \left\{ u(s, x) + \beta \int_{\mathbb{S}} V(s', \mu^*, p) \bar{Q}_{\mu^*}(ds' | s, x) \right\}. \quad (8)$$

Therefore, in line with Definition 5, the set of Berk-Nash equilibrium for a given primitive  $p$  is given by the set of fixed points of the equilibrium mapping  $T$ ,

$$\Lambda(p) \equiv \{(m^*, \mu^*) \in \mathcal{M}_1(\mathbb{S} \times \mathbb{X}) \times \mathcal{M}_1(\Theta) : z \in T(z, p)\}, \quad (9)$$

where  $T : Z \times P \rightarrow 2^Z$  is a set-valued function on the space of probability measures on states

---

<sup>22</sup>For finite state and action spaces, EP shows that convergence happens under a mild identification assumption. For infinite (but compact) state and action spaces, ADGK demonstrate that convergence occurs under identification only under the total-variation norm on the space of state and action distributions.

and actions and the set of parameters,  $Z = \mathcal{M}_1(\mathbb{S} \times \mathbb{X}) \times \mathcal{M}_1(\Theta)$  and  $P = \langle u, \beta, Q, Q_\Theta, \Theta \rangle$  are our primitives.<sup>23</sup>

For inquiring the comparative statics behavior of the equilibrium objects in  $\Lambda(p)$  with respect to the primitives of the environment, we need to define corresponding orders for the equilibrium objects and the primitives. For vectors  $x$  and  $y$  in  $\mathbb{R}^n$ , we use the following convention: “ $x \geq y$ ” means  $x_i \geq y_i$  in every component, “ $x > y$ ” means  $x \geq y$  and  $x \neq y$ , and “ $x \gg y$ ” means  $x_i > y_i$  in every component. A real-valued function,  $f : X \rightarrow \mathbb{R}$ , is said to be increasing if for  $x \geq y$  in the component-wise order,  $f(x) \geq f(y)$ , and convex if its domain is a convex set and for all  $x, y$  in its domain, and all  $\lambda \in [0, 1]$ , we have  $f(\lambda x + (1 - \lambda)y) \leq \lambda f(x) + (1 - \lambda)f(y)$ .<sup>24</sup> We follow Shaked and Shanthikumar (2007) in ranking probability measures on the state and action space, and the parameter space.

**Definition 6** (Shaked and Shanthikumar (2007)). *Let  $\mathcal{M}(X)$  denote the space of probability measures defined on a compact subset of  $X \subseteq \mathbb{R}^n$ .<sup>25</sup> For any two measures,  $\mu$  and  $\nu$  in  $\mathcal{M}(X)$ ,*

- (i)  $\mu$  usual order stochastically dominates  $\nu$ ,  $\mu \succeq_{st} \nu$ , if  $\int f(x)\mu(dx) \geq \int f(x)\nu(dx)$ , for any measurable, bounded, and increasing real-valued function  $f$ .
- (ii)  $\mu$  increasing and convex-order stochastically dominates  $\nu$ ,  $\mu \succeq_{icx} \nu$ , if  $\int f(x)\mu(dx) \geq \int f(x)\nu(dx)$ , for any measurable, bounded, increasing, and convex real-valued function  $f$ .

We rely on well-known orders for the primitives. For instance, the discount factor  $0 < \beta < 1$  is ranked in the natural order, as in higher patience implies a higher  $\beta$ . The true transition function  $Q$  can be ordered in multiple ways, including the usual stochastic order, the convex order, or the increasing convex order.<sup>26</sup> The parameter set  $\Theta$  is ranked in the strong-set order, where  $\Theta_2$  is greater than or equal to  $\Theta_1$  if, for any  $\theta$  in  $\Theta_2$  and any  $\theta'$  in  $\Theta_1$ , the maximum of  $\{\theta, \theta'\}$  belongs to  $\Theta_2$  and the minimum of  $\{\theta, \theta'\}$  belongs to  $\Theta_1$ . Changes in the primitives of the utility function refer to changes that impact utility levels such as variations in risk aversion or other factors, depending on the specific applications.

After defining all the necessary prerequisites for analyzing the comparative statics of the fixed points in Equation (9), the next logical step is to apply the standard lattice theoretic methods as used in (Hopenhayn and Prescott (1992), Topkis (1998)) and derive the monotonicity

<sup>23</sup>For the ease of the reader, the full expression for  $T$  is defined in the Proofs section, p.30.

<sup>24</sup>The interested reader is referred to Shaked and Shanthikumar (2007) for more details concerning univariate and multivariate orders. The results in this paper could potentially be extended to admit various orders.

<sup>25</sup>To be clear,  $X$  denotes a general set, not particularly the set of actions,  $\mathbb{X}$ .

<sup>26</sup>Convex order in this case, would refer to changes in true distribution in the sense of mean-preserving spreads. For two distributions,  $Q_1$  and  $Q_2$ , distribution  $Q_2 \succeq_{cx} Q_1$  iff  $\int f(s)Q_2(\cdot|s, x) \geq \int f(s)Q_1(\cdot|s, x)$  for every convex function,  $f$ , and for all  $(s, x) \in \mathbb{S} \times \mathbb{X}$ .

of the fixed points, the Berk-Nash equilibrium and the corresponding inferred model, in terms of the primitives. However, there is a technical challenge here. Notice that for our equilibrium mapping  $T$ , the underlying space on which it is defined is not a lattice. That is, the space of probability measures over states and actions (and parameters) are not lattices in any natural order.<sup>27</sup> Thus, standard lattice theoretic methods are of no use for our misspecified MDP setting. Therefore, to establish the monotonicity of the fixed points, our results rely on a rather novel application of a technical framework that originated in Smithson (1971) and was pioneered in the context of large dynamic economies by Acemoglu and Jensen (2015), for non-lattice spaces.

**Remark 1.** *Our setting fundamentally diverges from the setting in Acemoglu and Jensen (2015) in several key aspects. In their work, they examine an infinite-horizon dynamic economy featuring a continuum of agents subject to idiosyncratic shocks, whereas our setting deals with a single agent facing an infinite-horizon dynamic problem in which shocks are influenced by action choices. Most importantly, our analysis intertwines the agent’s decision-making process with learning under model misspecification, whereas in Acemoglu and Jensen (2015), the agents possess knowledge of the accurate distribution of idiosyncratic shocks, thus obviating the need for any learning component. This also furthers tailoring their techniques to our framework.*

Towards this end, we shall show (Lemmas 1 and 2 in Appendix B) that our equilibrium mapping  $T$  is monotonic in the sense defined below.

**Definition 7** (Smithson (1971)). Let  $X$  and  $Y$  be sets equipped with some partial order  $\succeq$ , and  $P$ , a partially ordered set. A correspondence  $T : X \times P \rightarrow 2^Y$  is Type I monotone in  $x$  for each  $p$  if for all  $x_1 \succeq x_2$  and  $y_2 \in T(x_2, p)$ , there exists  $y_1 \in T(x_1, p)$  such that  $y_1 \succeq y_2$ , and Type II monotone if for all  $x_1 \succeq x_2$  and  $y_1 \in T(x_1, p)$ , there exists  $y_2 \in T(x_2, p)$  such that  $y_1 \succeq y_2$ . A correspondence  $T : X \times P \rightarrow 2^Y$  is Type I monotone in  $p$  for each  $x$  if for all  $p_1 \succeq p_2$  and  $y_2 \in T(p_2, x)$ , there exists  $y_1 \in T(p_1, x)$  such that  $y_1 \succeq y_2$ , and Type II monotone if for all  $p_1 \succeq p_2$  and  $y_1 \in T(p_1, x)$ , there exists  $y_2 \in T(p_2, x)$  such that  $y_1 \succeq y_2$ .  $T$  is Type I (Type II) monotone if it is Type I (Type II) monotone in both  $X$  and  $P$ .<sup>28</sup>

By placing mild monotonicity structure on our environment, we shall show that the Berk-Nash equilibrium map  $T$  is Type I and Type II monotone. And then by appealing to an existence

<sup>27</sup>Let  $\mathcal{M}(X)$  denote the space of probability measures defined on a compact subset of  $X \subset \mathbb{R}^n$ . Although even if  $X$  is a lattice, the poset  $(X, \succeq_{st})$  is not a lattice as pointed in Kamae, Krengel, and O’Brien (1977). For e.g., for  $(\mathbb{R}^2, \succeq_{st})$  let  $p_1 = 0.5(\epsilon_a + \epsilon_b), p_2 = 0.5(\epsilon_a + \epsilon_c), p_3 = 0.5(\epsilon_c + \epsilon_b), p_4 = 0.5(\epsilon_a + \epsilon_d)$ , where  $a = (0, 0), b = (0, 1), c = (1, 0), d = (1, 1)$ . Then, both  $p_3$  and  $p_4$  are supremum, which is a contradiction.

<sup>28</sup>Smithson uses the term ‘multifunction’ in his paper. All the proofs in this section are given for Type I monotonicity. The proofs for Type II monotonicity follow analogously. It is important to point out here that while the statement for Theorem 1 is correct in Smithson (1971), its proof is wrong; the correct proofs are in Höft (1987).

result in Smithson (1971) for non-lattice spaces, we shall show that the set of fixed points is non-empty.<sup>29</sup> Finally, the following result in Acemoglu and Jensen (2015) will play a critical role in transfer the monotonicity of the equilibrium mapping  $T$  onto the set of its fixed points. We explain in more detail the applicability of these results in Section 4.

**Theorem** (Acemoglu-Jensen (2015)). *Let  $X$  be a compact topological space equipped with a closed partial order  $\succeq$ ,  $P$  a partially ordered set of primitives, and  $T : X \times P \rightarrow 2^X$  be upper hemicontinuous for each  $p \in P$ . Define the fixed-point correspondence,*

$$\Lambda(p) = \{x \in X : x \in T(x, p)\}.$$

*Then if  $T$  is Type I (Type II) monotone in  $p$ , so is  $\Lambda$  in  $p$ .*

### 3 Motivating Examples

We now present three examples of misspecified environments that illustrate the framework of this paper. Our first example is that of a misspecified AR(1) inference process where actions have no role. It is purely an inference problem.<sup>30</sup> Examples 2 and 3 are more substantive in that the choice of actions plays a role in learning, and vice-versa. They are based on EP and cover effort provision problems with misspecified effort-task dynamics and consumption-savings problem with misspecified wealth process, respectively. While Example 2 has an explicit analytical solution and is, therefore, used to outline the mechanics for comparative statics properties of a Berk-Nash equilibrium, the solution for Example 3 is intractable and, therefore, the most substantive for the applicability of this paper's techniques.<sup>31</sup>

**Example 1** (Inference of an AR(1) process). In this example, we characterize the Berk-Nash equilibrium of an inference process when the set of AR(1) models, parameterized by  $\theta$ , is misspecified. Let the state space  $\mathbb{S} = \mathbb{R}$ . Suppose the state variable  $s_{t+1}$  evolves via the true transition function,  $Q(\cdot|s_t)$ . Let the true process  $Q(\cdot|s_t)$  be an AR(1) process with parameter  $0 < |\rho| < 1$  and with the innovations distributed as a two-component mixture normal distribution with one component  $(\mu_1, \sigma_1^2)$  and the other component  $(\mu_2, \sigma_2^2)$ , and given by,

$$s_{t+1} = \rho s_t + \xi_{t+1}, \quad \xi_{t+1} \sim 0.5F_{(\mu_1, \sigma_1^2)} + 0.5F_{(\mu_2, \sigma_2^2)} \quad (10)$$

---

<sup>29</sup>The reader may wish to pause and review the material in the mathematical preliminaries in Appendix A.

<sup>30</sup>Per se, there is no decision process involved here. One could think of an economic situation where the flow of utility is constant over time, and therefore, actions play no role in either inference, nor, payoffs.

<sup>31</sup>We omit the full analysis of Examples 2 and 3 since they are already covered in the aforementioned papers.



where  $F$  denotes the cumulative density function for a normal distribution. The components have different means ( $\mu_1 \neq \mu_2$ ) but identical variances ( $\sigma_1^2 = \sigma_2^2$ ). An agent is equipped with a compact set  $\Theta$  of parameterized AR(1) models,  $Q_\theta(\cdot|s_t)$ , with Gaussian noise, and the process,

$$s_{t+1} = \theta s_t + \xi_{t+1}, \quad \xi_{t+1} \sim N(0, \sigma^2). \quad (11)$$




Figure 1: While the innovation process in the true distribution is mixture-normally distributed with distinct means and identical variances, resulting in a bi-modal distribution, the innovation processes in the agent’s models conform to a normal distribution.

The agent’s models are misspecified since the true distribution is not normally distributed, and therefore, is not in the set of models,  $Q_\Theta$ . Following Definition (5), the Berk-Nash equilibrium is the stationary distribution  $m_{\mathbb{S}}^*$  over the set of states (no actions), implied by the true process,  $Q(\cdot|s_t)$ , with the corresponding best-fit AR(1) parameter,  $\theta^* = \int_{\mathbb{S}} \hat{\theta}(s) m^*(ds)$ , where,

$$\hat{\theta}(s) \equiv \operatorname{argmin}_{\Theta} \operatorname{KL}(Q|Q_\theta) = \operatorname{argmin}_{\Theta} \mathbb{E}_Q \left[ \frac{\ln Q(\cdot|s)}{\ln Q_\theta(\cdot|s)} \right].$$

Now, given that the AR(1) models are Gaussian and therefore log-concave, the state-dependent best-fit model  $\hat{\theta}(s)$  is uniquely determined.<sup>32</sup> Since the innovations in the models are Gaussian,  $\hat{\theta}(s)$  is the least-squares minimizer with respect to the true transition  $Q(\cdot|s_t)$ , and solves the following,

$$\hat{\theta}(s) E_{Q(\cdot|s)}(s^2) = E_{Q(\cdot|s)}(s' s). \quad (12)$$

The best-fit inferred AR(1) parameter has the following form,

$$\theta^* = \int_{\mathbb{S}} \hat{\theta}(s) m_{\mathbb{S}}^* = \rho + \int_{\mathbb{S}} \frac{(\mu_1 + \mu_2)}{s} m_{\mathbb{S}}^*.$$

<sup>32</sup>It is worth mentioning here that Berk’s characterization of the limit posterior is for an i.i.d. process. Here conditioned on the state  $s$ , the AR(1) process behaves as i.i.d., and therefore,  $\hat{\theta}(s)$  could be interpreted as the posterior to which the agent converges if they were to receive infinite state-realizations, all i.i.d., given the state is fixed.

Now, let  $\mu_1 + \mu_2 = 0$ . That is, the mean of the components of the true distribution cancel each other. In this case, a Bayesian agent, starting from any prior on their models  $Q_\Theta$ , would eventually converge in the Berk-Nash equilibrium to the true persistence parameter  $\rho$  and yet continue to remain misspecified. For this case, the comparative statics of the steady state inferred  $\theta^*$  with respect to the AR(1) parameter of the true process is one-to-one in a trivial way.<sup>33</sup>

**Example 2** (Effort provision with unknown ability (EP)). In this setting of a dynamic effort provision problem, we follow EP in characterizing the Berk-Nash equilibrium. The equilibrium for this instance is a probabilistic prediction for the agent’s long run outcome to be either a success or a failure, along with their inference of their ability. This example is related to the growing literature on effort provision, as exemplified in Deimen and Wirtz (2022).

Each period, an agent chooses either to put high ( $H$ ) or low ( $L$ ) effort,  $x_t \in \mathbb{X} = \{L, H\}$ ,  $L < H$  in a task. The task then has two outcomes, either it fails (0) or succeeds (1), that is,  $s_{t+1} \in \mathbb{S} = \{0, 1\}$ . The payoffs are as follows,

$$u(x_t, s_{t+1}) = \begin{cases} s_{t+1}, & x_t = L, \\ s_{t+1} - c, & x_t = H. \end{cases}$$

In case the effort is low, the agent’s payoff is the state realization, while if the effort is high, in addition to the state realization, they also incur a cost  $c$  for their high effort. Under the true process  $Q(\cdot | s_t, x_t)$ , if the agent puts in the high effort, the probability of the task being a success is 1, irrespective of the outcome of the task in the current state. That is,  $Q(1 | s_t, H) = 1$ . With low effort, however, the probability of success tomorrow is dependent on the outcome today;  $Q(1 | 0, L) = q_0$ , if it is a failure, and  $Q(1 | 1, L) = q_1$ , if it is a success. Following EP, we assume that  $0 < q_0 < 1 - c < q_1 < 1$ . From a routine calculation in EP, the optimal correctly specified policy function is the following,

$$x(s_t) = \begin{cases} H, & s_t = 0, \\ L, & s_t = 1. \end{cases}$$

That is, a failure today, pushes the agent to work harder (since  $q_0 < 1 - c$ ), while a success makes them put lower effort (since  $1 - c < q_1$ ).

However, the agent has a set of misspecified models  $Q_\Theta = \{Q_\theta\}$ , parameterized by  $\theta \in \Theta = [0, 1]$ , where  $Q_\theta(1 | s, H) = 1$  and  $Q_\theta(1 | s, L) = \theta$  for all  $s \in \{0, 1\}$ . The agent’s models correctly captures the true dynamics, if they were to put in higher effort. However, under lower

---

<sup>33</sup>For the case when  $\mu_1 + \mu_2 \neq 0$ , the expression for  $\theta^*$  is potentially intractable since the expectation with respect to the Berk-Nash equilibrium  $m^*$  may simply not exist. The inspiration for the mixture normal distribution comes from a paper by Csaba and Szoke (2023) who look at learning when the set of likelihoods is misspecified.

effort, their models imply that the chances of success are independent of the state today, and is denoted by the parameter  $\theta$ . Following a growing literature on effort provision that models an agent's control over their future success, based on their effort and ability, we novelly interpret the parameter  $\theta$  as an index for the agent's ability level. That is, the agent is someone who thinks that when the effort is high, then their ability plays no role in determining the outcome tomorrow since higher effort guarantees success. Whereas, when the effort is low, it is only the ability that matters for the successful outcome, with higher ability implying higher chances of success.

For the Berk-Nash equilibrium, following EP, we shall focus on the unique mixed-strategy equilibrium wherein the choice of actions are independent of the current state, as implied by the misspecified models. A routine calculation shows that under certain parameterizations, a unique Berk-Nash equilibrium exists wherein the agent plays a mixed strategy, conditioned on the state,  $m^*(L|0) = m^*(L|1) = \frac{q_1 - (1 - c)}{c(q_1 - q_0)}$ , and correspondingly infers their ability,  $\theta^* = 1 - c$ .<sup>34</sup>

**Example 3** (Savings with misperceived wealth process (EP, ADGK)). In this example, we demonstrate that there are environments wherein there is no clear analytical solution for the Berk-Nash equilibrium and its corresponding inferred model, a feature most common to dynamic programming problems.

In each period, an agent realizes wealth  $y_t$ , an i.i.d. preference shock  $z_t$ , and subsequently chooses to save  $x_t \in [0, y_t] = \mathbb{X} \subseteq \mathbb{R}_+$ . The utility in period  $t$  is given by  $u(y_t, z_t, x_t) = z_t \ln(y_t - x_t)$ . The state variables, denoted as  $s = (y, z)$ , belong to  $\mathbb{S} = \mathbb{R}_+ \times [0, 1]$ . The evolution of wealth  $y_{t+1}$  in the following period is governed by the following process,

$$\ln y_{t+1} = \alpha^* + \beta^* \ln x_t + \varepsilon_t,$$

where the unobserved productivity shock,  $\varepsilon_t = \gamma^* z_t + \xi_t$ , with  $\xi_t$  following a standard normal distribution,  $\xi_t \sim \mathcal{N}(0, 1)$ , and  $z_t$  following a uniform distribution, with mean  $E(z)$ . Here,  $\beta^*$  denotes the true return on one additional unit of log of saving in terms of the log of wealth. We assume that it lies between 0 and 1, so that the true Markov process is stationary. Following EP, we shall assume that the preference and the productivity shocks are positively correlated,  $\gamma^* > 0$ . For an agent that knows all the primitives of the environment, solving for the Bellman as in Equation (3), the correctly specified optimal policy function is given by,  $x^* = A_z(\beta^*)y$ , where  $A_z(\beta^*) = \frac{\delta\beta^*E(z)}{(1 - \delta\beta^*)z + \delta\beta^*E(z)}$ , where  $0 < \delta < 1$  is the discount factor.

---

<sup>34</sup>The steady-state marginal distribution is given by  $m_S^*(1) = \frac{1 - c - q_0}{q_1 - q_0}$ , and  $m_S^*(0) = 1 - m_S(1)$ . Further,  $\theta_Q(m) = (1 - m_S(1))q_0 + m_S(1)q_1$ . The reader is referred to the EP paper (pp. 726-728) for the full numerical details.

However, the agent's set of models postulates the following wealth process,  $\ln y_{t+1} = \alpha + \beta \ln x_t + \varepsilon_t$ , where  $\varepsilon_t \sim N(0, 1)$ , thus ignoring the correlation between the productivity and the preference shocks. So while the true process is  $Q(y', z' | y, z, x)$  is such that  $y'$  and  $z'$  are independent,  $y'$  has a log-normal distribution with mean  $\alpha^* + \beta^* \ln x + \gamma^* z$  and unit variance, and  $z' \sim U[0, 1]$ , the misspecified process is  $Q_\theta(y', z' | y, z, x)$  is such that  $y'$  and  $z'$  are independent,  $y'$  has a log-normal distribution with mean  $\alpha + \beta \ln x$  and unit variance, and  $z' \sim U[0, 1]$ .

For this case, the Berk-Nash equilibrium is characterized by an optimal policy function,  $x^m = A_z(\beta^m) y = \frac{0.5\delta\beta^m}{(1 - \delta\beta^m)z + 0.5\delta\beta^m} y$ , where there exists a corresponding inferred model  $\beta^m \in (0, \beta^*)$ . That is, the misspecified agent underestimates the return on her saving,  $\beta^m < \beta^*$ , and therefore, undersaves in the Berk-Nash equilibrium,  $x^m < x^*$ .<sup>35</sup> Notice however that there is no closed-form expression for  $\beta^m$ , that outlines its dependence on the environment primitives, and therefore, how the inferred return, the equilibrium policy function, and the invariant distribution respond to the primitives of the environment is an open question.<sup>36</sup>

## 4 Main Results

We now turn to our main results that are on the existence and the comparative statics of Berk-Nash equilibrium for a misspecified dynamic optimization problem. We invoke techniques from the fixed point literature (Smithson (1971); Acemoglu and Jensen (2015)) and, in doing so, impose further structure on our regular SMDP. While assumptions 1 and 2 are useful for establishing increasing policy selections, assumptions 3 and 4 give the required monotonicity and identification properties for our models. Throughout, we shall assume that  $Q$  is monotone. The proofs for all the results are in Appendix B.

**Assumption 1.** *The state, action, and parameter spaces are lattices, and for any primitive  $p \in P$ ,  $u(s, x, p)$  is supermodular in  $(s, x)$  and increasing in  $s$ .*

Our first assumption is rather standard and emphasises a lattice structure on the state,

---

<sup>35</sup>The intuition behind this is that a higher preference shock  $z$  is associated with a lower saving proportion,  $A_z(\beta^m)$ , since it leads to a lower inferred belief of return on saving,  $\beta^m$ , since the preference and productivity shocks are positively correlated.

<sup>36</sup>From EP: The inferred model at the stationary Berk-Nash equilibrium solves the following equation,

$$\hat{\beta}(A_z(\beta)) = \beta^* + \gamma^* \frac{\text{Cov}(z, \ln A_z(\beta))}{\text{Var}(\ln A_z(\beta)) + \text{Var}(\ln Y)}, \quad (13)$$

where the covariance and variance are taken with respect to the true distribution,  $Q(y', z' | y, z)$ . Under the parametric assumptions of  $\gamma^* > 0$ , and  $\text{Cov}(z, A_z) < 0$ , EP establish that there exists a  $\beta^m$ , the inferred model in the Berk-Nash equilibrium. Further, in EP, notice that the inferred model depends on the proportion,  $A_z$ , instead of the distribution over states and actions, which is unlike the case required as per the conditions of Berk-Nash equilibria.

action, and parameter spaces. As evinced by the examples in this paper and in general, for many economic applications, the state, action, and parameter spaces are a subset of the real line  $\mathbb{R}$ , for which the lattice assumption is trivially satisfied in the natural order. Further, we assume that the payoff function is supermodular in state and action and is increasing in the state variable. A payoff function  $u : \mathbb{S} \times \mathbb{X} \rightarrow \mathbb{R}$  is supermodular if, for all  $(s_1, x_1), (s_2, x_2) \in \mathbb{S} \times \mathbb{X}$ ,  $u((s_1, x_1) \vee (s_2, x_2)) + u((s_1, x_1) \wedge (s_2, x_2)) \geq u(s_1, x_1) + u(s_2, x_2)$ . This implies that states and actions exhibit complementarity in the sense that marginal contribution to the payoff of increasing the action increases with a higher state. The supermodularity in states and actions and monotonicity (in state) of the payoff function alongwith our next assumption guarantees that optimal policy correspondence is increasing in the state variable, for any given model distribution  $\mu$ , and primitive  $p \in P$ .

**Assumption 2.** *The following holds true for all models in the family of models,  $\mathcal{Q}_\Theta = \{Q_\theta : \theta \in \Theta\}$ . For any increasing real-valued function  $f(\cdot)$ ,*

- (i)  $Q_\theta$  is stochastically increasing in  $(s, x)$  i.e.  $\int_{\mathbb{S}} f(s') Q_\theta(ds'|s, x)$  is increasing in  $(s, x)$ .
- (ii)  $Q_\theta$  is stochastically supermodular in  $(s, x)$  i.e.  $\int_{\mathbb{S}} f(s') Q_\theta(ds'|s, x)$  is supermodular in  $(s, x)$ .

Notice that this assumption is exclusively on the set of models. It requires the model transition functions to be stochastically increasing and stochastically supermodular in states and actions. This implies that for every model  $Q_\theta$ , a higher current state and action increases the probability of observing a higher state in the next period and that an incremental amount of action increases this probability, the lower the state is in the current period.<sup>37</sup> Assumptions 1 and 2 are consistent with the structure one requires for a correctly specified environment and the reader is referred to Theorem 3.9.2 in Topkis (1998) for a similar set of assumptions. For a given model distribution  $\mu$ , and primitive  $p$ , the above two assumptions guarantee an increasing optimal policy correspondence  $G(s, \mu, p)$ .

We next assume that given any observable endogenous data, i.e. given any distribution  $m$  over states and actions, generated by agent's endogenous learning and decision making process as in Equation (1), the best-fit parameter is point-identified. That is, given  $m$ , the best-fit set is singleton and therefore, the parameter is uniquely determined.

**Assumption 3.** *For any given  $m \in \mathcal{M}_1(\mathbb{S} \times \mathbb{X})$ , a SMDP  $(Q, Q_\Theta)$  is point-identified if  $\theta, \theta' \in \Theta(m; Q)$  implies that  $\theta = \theta'$ . It is point-identified if it is point-identified for all  $m \in \mathcal{M}_1(\mathbb{S} \times \mathbb{X})$ .*<sup>38</sup>

<sup>37</sup>For example, consider the following AR(1) process,  $s_{t+1} = \theta s_t + \epsilon_{t+1}$ , where  $\epsilon_{t+1}$  is distributed normally with mean 0 and variance  $\sigma^2$ . Then for every  $\theta$ ,  $Q_\theta$  is stochastically increasing. See Examples 2 and 3 for illustrations on stochastic supermodularity.

<sup>38</sup>The reader is referred here to Lewbel (2019) for an exposition on identification problem in economics.

The identification assumption is satisfied for many applications in the existing misspecification literature and also plays an important role in the convergence of the steady state behavior of an agent in a SMDP to the Berk-Nash equilibrium; see examples in Esponda and Pouzo (2016), EP, and, Esponda, Pouzo, and Yamamoto (2021). For our comparative statics purposes, identification ensures that we always get a unique selection  $\theta_Q(m)$  from the best-fit set  $\Theta(m; Q)$ . A function  $\theta_Q(m)$ , maps a distribution  $m$  into a real number in  $\Theta$  is said to be increasing in the usual stochastic order  $\succeq_{st}$ , if  $\theta_Q(m_2) \geq \theta_Q(m_1)$  whenever  $m_2 \succeq_{st} m_1$ . For our last assumption, we require the weighted Kullback-Liebler divergence to exhibit the single-crossing property in the parameter  $\theta$  and distribution  $m$ .

**Assumption 4.**  $K_Q(\theta; m)$  satisfies the single crossing property in  $(\theta; m)$ , that is, for  $\theta_1 \leq \theta_2$ , and for  $m_2 \succeq_{st} m_1$ ,  $\delta(m_1) = K_Q(\theta_2, m_1) - K_Q(\theta_1, m_1) \geq (<)0$  implies  $\delta(m_2) = K_Q(\theta_2, m_2) - K_Q(\theta_1, m_2) \geq (<)0$ .

The assumption above is necessary and sufficient to imply that the best-fit set  $\Theta(m; Q)$  (coupled with Assumption 3, the inference function  $\theta_Q(m)$ ), is increasing in the strong-set order on the real line.<sup>39</sup> However, as is well-known in the comparative statics literature, it can be difficult to verify for single-crossing differences. Therefore, a sufficient and easy to check condition to guarantee an increasing best-fit function or set is to require that the models follow a *expected* log-likelihood property.

**Sufficient Condition 1:** For any two models  $\theta_1, \theta_2 \in \Theta$ , such that  $\theta_1 < \theta_2$ , define the log-likelihood ratio,  $L(s'|s, x) = \ln(D_{\theta_2}(s'|s, x)) - \ln(D_{\theta_1}(s'|s, x))$ , where  $D_{\theta}(s'|s, x)$  are the Radon-Nikodym derivatives, introduced for measuring KL divergences in (5). Then the models are said to follow the generalized likelihood ratio property if the expectation of  $L$  with respect to the true distribution  $Q(\cdot|s, x)$  is increasing in state and actions,  $(s, x)$ .<sup>40</sup>

We can now state our first result on the existence of a Berk-Nash equilibrium for a regular SMDP with infinite states and actions.

**Theorem 1.** *Under assumptions 1-3, every regular SMDP  $(Q, \mathcal{Q}_{\Theta})$  with a bounded and continuous utility function has a Berk-Nash equilibrium and the set of such equilibria is compact.*

EP and ADGK provide proofs for the existence of a Berk-Nash equilibrium in an SMDP with finite and infinite environments, respectively. While the former proves the theorem for

---

<sup>39</sup>This is similar to assuming *monotone* Bayesian updating of the kind in Torres (2005) where stochastically higher states lead to higher posterior probabilities of the parameter. Also, since the parameter spaces are in  $\mathbb{R}$ , quasi-supermodularity is trivially satisfied; see Milgrom and Shannon (1994).

<sup>40</sup>Ross (1987) contains an useful discussion regarding the existence of the expectation of the log-likelihood ratio statistic. For instance, if the true, and the model distributions belong to the Gaussian family, this statistic would always exist.

finite states and actions, ADGK takes the finite result of EP as given and extends it to more naturally appealing instances of infinite environments, using novel tools in non-standard analysis. In contrast, our existence result relies on the assumed monotonicity and identification properties of the structure of our problem. It is important to mention here that we do rely on the regular SMDP definition in ADGK, most importantly on part (iv) of Definition 3, because it allows us to consider distributions that have unbounded Radon-Nikodym derivatives. This is particularly relevant when dealing with distributions over infinite spaces, such as the AR(1) process.<sup>41</sup> For unbounded state spaces, one continues to rely on ADGK's structure and existence results. Further, the set of Berk-Nash equilibria will be compact and in particular, there will always exist the least and the greatest Berk-Nash equilibrium,  $m^*$  and corresponding beliefs  $\mu^*$  over the best-fit inferred models.<sup>42</sup>

Following Acemoglu and Jensen (2015), we next define a positive shock for a SMDP. Notice that this positive shock is defined for any  $p \in P = \langle u, \delta, Q, Q_\Theta, \Theta \rangle$  and therefore, appeals to any primitive  $p$  that can be considered.

**Definition 8.** For any given belief  $\mu$ , a change in a primitive of the SMDP from  $p_1$  to  $p_2$  is a *positive (negative)* shock if  $G(s, \mu, p_2)$  is greater than  $G(s, \mu, p_1)$  in the strong set order. That is, for all  $y_1 \in G(s, \mu, p_1)$  and  $y_2 \in G(s, \mu, p_2)$ , the join  $y_1 \vee y_2 \in G(s, \mu, p_2)$  and the meet  $y_1 \wedge y_2 \in G(s, \mu, p_1)$ . Further, for a given primitive  $p \in P$ , a change in beliefs over models from  $\mu_1$  to  $\mu_2$  is a *positive (negative)* shock if  $G(s, \mu, p)$  is ascending in  $\mu$  from  $\mu_1$  to  $\mu_2$ , that is, if  $y_1 \vee y_2 \in G(s, \mu_2, p)$  and  $y_1 \wedge y_2 \in G(s, \mu_1, p)$  for all  $y_1 \in G(s, \mu_1, p)$  and  $y_2 \in G(s, \mu_2, p)$ .

We now state the main result of this paper.

**Theorem 2.** *Suppose assumptions 1-4 hold. Then a positive shock to the primitives of the regular SMDP will lead to an increase in the least and the greatest equilibrium best-fit models.<sup>43</sup> Further, a positive shock to the primitives will lead to*

- (a) *an increase in the least and greatest Berk-Nash equilibrium in the usual stochastic order dominance if changes in beliefs over models are positive shocks, and*
- (b) *a decrease in the least and greatest Berk-Nash equilibrium in the usual stochastic order dominance if changes in beliefs over models are negative shocks.*

Theorem 2 highlights the two-way interaction between decision making and learning for monotone comparative statics behavior in misspecified dynamic optimization problems with learning.

---

<sup>41</sup>Although our examples are based on unbounded state spaces, our results in this paper are established for compact state and action spaces.

<sup>42</sup>The argument for the least and greatest follows from Theorem 4 in Acemoglu and Jensen (2015) for the non-lattice case. Also, see Footnote 9 (pp. 1389) in Hopenhayn and Prescott (1992) for several instances of compact subsets of measures in economic problems.

<sup>43</sup>It is important to note here that we are assuming that the constraint on the parameter set is non-binding.

Under assumptions 1-4 on the regular SMDP, it predicts that a positive shock to any of its primitives will lead to an increase in the least and greatest equilibrium beliefs over the best-fit parameterized models. Furthermore, a positive (negative) shock will lead to an increase in the least and the greatest Berk-Nash equilibrium, the steady state distribution over states and actions, if changes in beliefs over models is a positive (negative) shock. For instance, if there is a unique Berk-Nash equilibrium, as in the examples here and in the literature, with a corresponding inferred model, then Theorem 2 gives a clear prediction for the comparative statics of the equilibrium objects.

**Remark 2.** *Our framework is more general to that of Heidhues, Kőszegi, and Strack (2018), and therefore, the sharp predictions that they have regarding “perverse” comparative statics for a stylized production problem with a loss function is something that we are not able to explain here.*

While the full proof of the main result is provided in the Appendix, we provide a brief proof sketch below that outlines the main points of our argument.

**Short proof sketch:** The proof of Theorem 2 follows a three-step proof structure as in Acemoglu and Jensen (2015). The first step involves showing that under Assumptions 1 and 2, and for any fixed model distribution  $\mu$ , the set of Berk-Nash equilibrium  $m$  obtained via the fixed points of the equilibrium mapping  $T$  will be Type I and Type II increasing in the primitives  $p$ . Using the set of fixed points in the first step, the second step involves constructing a mapping  $\hat{\theta}$  that for any model distribution  $\mu$  and primitive  $p$ , gives a set of model distributions. This mapping is constructed from part (b) of Definition 5 wherein for each of the Berk-Nash equilibrium  $m$  obtained in step 1, the construction gives a set of  $\mu$ 's. It is the fixed points of this map that are our equilibrium distribution given  $p$ . The *third* and *final* step shows that if the least and the greatest selections of the mapping  $\hat{\theta}$  are increasing in primitives, then the associated fixed points are increasing in  $p$ .

**Remark 3.** Acemoglu and Jensen (2015) gives several sufficient conditions to identify positive shocks for their environment. All of their results (Lemmas 1-3) translate to our case,<sup>44</sup> albeit with some moderation for the endogenous dynamic programs. For example, a change in the primitive  $p$  that influences the decision problem through the utility function, such as the risk-aversion parameter, is a positive shock if the utility function  $u(s, x, p)$  has increasing differences in  $x$  and  $p$ . Similarly, under Assumptions 1 and 2, a change in discount factor  $\beta$  and model distribution  $\mu$  is a positive shock.

Our next theorem is on the inference of a Bayesian learner in a misspecified MDP who undergoes an *expansion* of their set of models in the strong-set order.

---

<sup>44</sup>The reader is referred to pp. 604-605 in Acemoglu and Jensen (2015).



**Theorem 3.** *Suppose the hypothesis in Theorem 1 continue to hold. If a change in beliefs over models is a positive shock, then an increase in the parameter set under the strong set order leads to an increase in the least and the greatest equilibrium best-fit models.*

Our last assumption modifies the single crossing property for instances when distributions over states and actions are ordered in the increasing and convex stochastic order.

**Assumption 5.**  $K_Q(\theta; m)$  satisfies the single crossing property in  $(\theta; m)$ , for increasing convex order  $\succeq_{icx}$  if, for  $\theta_1 \leq \theta_2$ , and for  $m_2 \succeq_{icx} m_1$ ,  $\delta(m_1) = K_Q(\theta_2, m_1) - K_Q(\theta_1, m_1) \geq (<)0$  implies  $\delta(m_2) = K_Q(\theta_2, m_2) - K_Q(\theta_1, m_2) \geq (<)0$ .

**Sufficient Condition 2:** For any two models  $\theta_1, \theta_2 \in \Theta$ , such that  $\theta_1 < \theta_2$ , define the log-likelihood ratio,  $L(s'|s, x) = \ln(D_{\theta_2}(s'|s, x)) - \ln(D_{\theta_1}(s'|s, x))$ . Then the models are said to follow the expected likelihood ratio property in increasing and convex order if the expectation of  $L$  with respect to the true distribution  $Q$  is increasing and convex in state and actions,  $(s, x)$ . Our final theorem addresses the monotonicity of Berk-Nash equilibrium for the increasing and convex order.

**Theorem 4.** *Suppose assumptions 1-3 and assumption 5 holds. Then a positive (negative) shock to the primitives will lead to an increase in the least and greatest Berk-Nash equilibrium in the increasing convex order if increases in beliefs over models are positive (negative) shocks.*

## 5 Analysis of Examples

In this section, we show the applicability of our results to Examples 2 and 3, where learning and decision-making are interlinked. Example 2 has an explicit analytical solution and is more of an illustrative example of the structure employed in this paper. However, Example 3, like most problems in dynamic programming, does not have a closed-form solution and, therefore, is most amenable to analysis with our framework.

**Example 2 (contd.)** (Dynamic effort with unknown ability (Esponda and Pouzo (2021))). We now verify our assumptions. The state space  $\mathbb{S} = \{0, 1\}$  and action space  $\mathbb{X} = \{H, L\}$  are lattices. Further, the utility function is increasing in the state variable  $s_{t+1}$  and for any given cost of effort  $c$ , it is also supermodular in state and actions owing to its linearity. Therefore, Assumption 1 is satisfied. To check for Assumption 2, the expectation of any increasing function with respect to model transition functions should be increasing and supermodular in states and actions. This is again satisfied given the structure of the transition functions since, irrespective of the state, higher action lead to success with probability 1 in the next period, and given the

action, the probability of success is independent of the current state. Since the weighted KL-divergence is strictly concave, for every given  $m > 0$  we have a unique minimizer, and hence, it is identified and satisfies Assumption 3. Therefore, a Berk-Nash equilibrium exists for this SMDP from Theorem 1. Notice that Assumption 4 is satisfied since the mapping  $\theta_Q(m_s)$  is increasing in  $m_s$ . Therefore, our results in Theorem 2 are applicable for this setting. We focus on the case with unique solution and for this example, the misspecified policy function is independent of the current state. That is,

$$g(s, \theta, p) = \begin{cases} H, & \theta < 1 - c, \\ L, & \theta \geq 1 - c. \end{cases}$$




Figure 2: An increase in the cost  $c$  of high effort is a negative shock for the optimal policy correspondence, which in this case is a function.

We now verify the prediction of our results vis-à-vis the analytical solution. A fall in the cost of effort  $c$  is a positive shock, and therefore, leads to a higher inference of ability,  $\theta^* = 1 - c$ , in the Berk-Nash equilibrium, as predicted in Theorem 2. Similarly, a fall in the cost of the effort leads to an increase in the probability of success at the steady state  $m_s^*(1)$ ; since an increase in the model parameter is a positive shock. Here, the stationary distribution  $m_s^*$  is increasing in the usual (first) stochastic order dominance in the cost of effort.

**Example 3 (contd.)** (Savings with misperceived wealth process (EP, ADGK)). First, we verify our assumptions. The state, action and parameter spaces are lattices and the utility function is increasing in the state variables,  $y$  and  $z$ . Since the payoff function is concave in  $y$  and  $x$ , and  $\frac{d^2u(y, z, x)}{dx dy} > 0$ , it is supermodular, and hence, satisfies Assumption 1. Further, the model distributions are Gaussian with mean  $\alpha + \beta \ln x$ , and unit variance and therefore, satisfy Assumption 2. This follows from the fact that an increase in action  $x$ , increases the

mean, and shifts the model distributions rightward, in the usual order stochastic dominant sense. Assumption 3 is met since the Gaussian distribution is strictly log-concave, and therefore, ensures unique identification. Hence, from Theorem 1, the Berk-Nash equilibrium exists. Furthermore, Assumption 4 is satisfied since this follows from a routine verification of the sufficient condition 1 for increasing best-fit models, under the usual stochastic order.<sup>45</sup>

Therefore, from Theorem 2, we have a clear prediction about the steady-state inferred model and the Berk-Nash equilibrium. For example, an increase in the discount factor, a positive shock to the primitives, will lead to a higher inferred return  $\beta^m$  in the equilibrium. Similarly, an increase in the expected preference shock, another instance of a positive shock, will also lead to a higher inferred return. Moreover, since increases in beliefs about the return are a positive shock, it will perpetuate into greater savings in the steady state, and a usual order stochastic increase in the stationary wealth distribution. One advantage of our comparative statics framework is that we generate predictions about the comparative statics of the equilibria even when a closed-form solution may not exist for the equilibrium objects, as is true for this case. Further, since Assumption 5 holds, the consequent implications also hold for Theorem 4.

## 6 Welfare Comparisons

We now turn towards comparing the *objective* welfare of an agent facing a regular SMDP, under correctly specified and misspecified Bayesian learning and characterize an upper bound on the difference between the two instances, the costs of misspecification, in terms of the primitives of the environment. For the purposes of this section, we shall assume that the utility function is bounded and strictly concave. Further, we shall assume that  $u : \mathbb{S} \times \mathbb{X} \rightarrow \mathbb{R}$  is continuously differentiable in actions.

Let  $g(s, \bar{\theta})$ , parameterized by a parameter  $\bar{\theta}$ , denote an optimal policy function. Suppose it exists under both correctly specified and misspecified environments, where  $\bar{\theta}$  take values  $\theta^*$  and  $\theta_*$  under correctly specified and misspecified learning, respectively. We are interested in comparing the welfare of a correctly specified agent with parameter  $\theta^*$  to the welfare of an agent who settles at the parameter  $\theta_*$  asymptotically, under misspecified Bayesian learning as in Equation (1). Then, the agent's welfare under the parameter  $\bar{\theta}$ , is the *objective* ex-ante expected discounted payoff,  $W(s, \bar{\theta})$  of choosing the optimal action,  $g(s, \bar{\theta})$ , and is given as,

$$W(s, \bar{\theta}) = \mathbb{E}_{Q(\cdot|s, g(s, \bar{\theta}))} \left[ \sum_{t=0}^{\infty} \beta^t u(s_t, g(s_t, \bar{\theta})) \right], \quad t = 0, 1, 2, \dots \quad (14)$$

---

<sup>45</sup>The expression  $\mathbb{E}_{Q(\cdot|s, x)} L(s'|s, x)$  is increasing in  $(s, x)$  under the Gaussian distributions,  $Q$  and  $Q_\theta$ . For  $(\alpha_1, \beta_1) < (\alpha_2, \beta_2)$ ,  $\mathbb{E}_{Q(\cdot|s, x)} L(s'|s, x) = \frac{1}{2\sigma^2} [(\alpha_2 - \alpha_1)(\beta_2 - \beta_1)(\gamma^* z)x]$  which is increasing in  $x$  and  $z$ .

Notice that for objective welfare, the expectation is solely with respect to the true transition function,  $Q(\cdot|s, g(s, \bar{\theta}))$ . In Equation (14), the form of environment affects welfare through the optimal policy action via two channels: first, through the per-period utility function, and second, through the true transition function,  $Q$ . The optimal policy action  $g(s, \theta_*)$  under misspecified Bayesian learning can theoretically be computed in accordance with Definition 5 and the *subjective* distribution  $Q_{\theta_*}$ . Our next theorem compares the welfare ranking between the two instances where our choice of metric is the one induced by sup-norm in the space of functions,  $W : \mathbb{S} \times \Theta \rightarrow \mathbb{R}$ .<sup>46</sup>

**Theorem 5.** *Welfare under correctly specified learning  $W(s, \theta^*)$  is weakly greater than welfare under misspecified learning,  $W(s, \theta_*)$ . Further, if  $m_0$  and  $m_1$  denote the absolute upper bound on the utility function and the marginal utility function, respectively, and if  $\|g(s, \theta^*) - g(s, \theta_*)\| \leq \gamma$ , then*

$$\|W(s, \theta^*) - W(s, \theta_*)\| \leq \frac{2\beta m_0(1 - e^{-k^*}) + m_1\gamma}{1 - \beta}, \quad (15)$$

where  $k^*$  is the upper bound on the KL divergence of  $Q(\cdot|s, g(s, \theta_*))$  with respect to  $Q(\cdot|s, g(s, \theta^*))$ .

We interpret the difference  $\gamma$  between the two policy functions in the sup-norm as the approximation error in the space of policy functions. Then, for a given approximation error  $\gamma$ , Equation (15) supplies an upper bound on the welfare comparison in terms of the primitives, namely the discount factor, and the absolute bounds on the utility, the marginal utility function, and the KL divergence. Notice that if  $\gamma = 0$ , that is, the policy functions under the two instances of learning are identical, then the two welfare quantities are equal, by construction. The upper bound in Equation (15) gives intuitive comparative statics, vis-a-vis the model primitives. For instance, the discrepancy in welfare is larger, the greater the approximation error,  $\gamma$ . Similarly, for a positive approximation error, the discrepancy increases in the discount factor. That is, as the agent gets increasingly patient, reflected in a higher discount factor  $\beta$ , his approximation error keeps accumulating over the horizon, leading to a greater discrepancy in welfare. Furthermore, higher bounds on the utility function and its corresponding marginal utility function lead to a greater discrepancy in the welfare. The marginal utility function plays a role in the discrepancy by attaching itself to the approximation error. Intuitively, getting the policy wrong is more pronounced if it matters at the margins. Notice that for the upper bound on the KL divergence, we only compare the implied distributions by the true transition function  $Q$ , under the two policy functions. This is because our interest lies in comparing the objective welfare under  $Q$ .

It is worth mentioning here that bound in Equation (15) is not specific to a misspecified MDP, rather, it holds for any MDP environment with a endogenous state evolution process,

---

<sup>46</sup>The metric induced by sup-norm is the following:  $\|W(s, \theta^*) - W(s, \theta_*)\| = \sup_s |W(s, \theta^*) - W(s, \theta_*)|$ . Similarly, for the policy function  $\|g(s, \theta^*) - g(s, \theta_*)\| = \sup_s |g(s, \theta^*) - g(s, \theta_*)|$ .

when plugged with alternate policy functions. In this sense, it is related to Lemma 3.1 in Santos (2000) which does a similar exercise in bounding welfare discrepancies. However, the state process in there is exogeneous. Furthermore, Santos (2000) uses Euler residuals to provide upper bounds on the approximation error, in terms of the primitives, for both the policy and the value functions, which is useful in the context of computing numerical approximations. Often, MDP environments do not have closed form solutions, thus making economists rely on numerical techniques to solve for them computationally. We hope that Theorem 5 could be instructive in this regard, and help in the numerical approximation of the Berk-Nash equilibria.

## 7 Concluding Remarks and Extensions

Models, by their very nature, offer simplified abstractions of reality, inevitably omitting certain nuances. This paper delineates conditions under which an important qualitative property, the comparative statics of decision-making and the corresponding inference, is preserved under model misspecification. The main contribution of this paper lies in establishing monotone comparative statics results for misspecified dynamic optimization problems, by a novel application of techniques in the fixed points literature (Smithson (1971)) and used in the context of large dynamic economies (Acemoglu and Jensen (2015)). We also illustrate the utility of these results for general interest environments. Further, we provide an upper bound on the welfare between correct and misspecified learning, in terms of the primitives. We conclude by outlining several promising directions for extending our findings to more broader environments.

**Multi-dimensional parameter spaces.** Since most of the current applications in the literature predominantly involve one-dimensional models,<sup>47</sup> the results presented in the paper assume the model (parameter) set to be a compact subset of the real line. However, they could be extended to multi-dimensional parameter spaces. This would involve a suitable rehabilitation of Assumptions 3 and 4 for lattice and non-lattice multi-dimensional parameter spaces. Although, one limitation of such a result using the current techniques is that the comparative statics results imply that all inferred parameters in the equilibrium must be monotonic together, which may not be ideal for specific economic applications.

**Other forms of updating.** While our results primarily address the asymptotic behavior of Bayesian learners, it's important to note that other inference processes, such as Maximum Likelihood Estimation (MLE) and moment-based learning, can lead to the same asymptotic inferences. Csaba and Szoke (2023) provide an illustrative example that demonstrates the identical nature of inference between Bayesian inference and methods based on likelihoods.

---

<sup>47</sup>This feature of the current literature is also noted in Esponda, Pouzo, and Yamamoto (2021).

**Static environments and dynamic concerns.** Our results are tailored for dynamic MDP environments but they potentially are also applicable, subject to modifications, to static environments, such as the one in Esponda and Pouzo (2016) for static games. This also spills over to settings where agents are endogenously concerned about their misspecified models, such as those in Lanzani (2022).<sup>48</sup>

**Monotone comparative dynamics.** The scope of our paper is limited to analyzing the comparative statics properties of the equilibrium objects in the steady state. An interesting question arises regarding how the dynamics of the misspecified learning process might react to changes in its primitives. Specifically, we do not know how the variations in the frequency of state-action pairs, as well as the sequence of posteriors as described in Equation (1), would be influenced because of changes in the economic primitives. There is added complexity in the dynamics due to potential complementarities between current actions and the inference process via the role of experimentation. In this regard, the tools developed in Balbus, Dziewulski, Reffett, and Woźny (2022) for monotone comparative dynamics for stochastic games could be useful for future exploration.

**Misspecification in large economies.** Our results are also applicable to misspecified dynamic economies with a continuum of agents, such as the one conceived in Molavi (2019). Given that the techniques we rely on (Acemoglu and Jensen (2015)) were developed in the context of economies with a continuum of agents, the results in this paper should potentially be straightforward to apply on equilibrium concepts, Molavi (2019) develops for the boundedly rational agents working with misspecified models in macroeconomic environments.

---

<sup>48</sup>The reader is referred to Corollary 1 in Lanzani (2022).

## A Mathematical Preliminaries

In this section, we provide the necessary mathematical preliminaries required to go through the proofs. Let  $X$  be a set. A subset  $\succsim$  of  $X \times X$  denotes a *binary relation* on  $X$ . A binary relation  $\succsim$  is a partial order if it is reflexive, transitive, and anti-symmetric. A partially ordered set, or a *poset*, is a pair  $(X, \succsim)$  that consists of a set  $X$  and a partial order  $\succsim$ . A binary relation  $\succsim$  is closed if the graph of  $\succsim$  is a closed subset of  $X \times X$ . The poset  $(X, \succsim)$  is a lattice if for any  $x, x' \in X$ , the greatest lower bound (infimum)  $x \wedge x'$  and the least upper bound (supremum)  $x \vee x'$  are in  $X$ , where  $\wedge$  and  $\vee$  denote the meet and join operations, respectively. A subset  $A \subseteq X$  is a sublattice of lattice  $X$  if  $A$  is a lattice that contains the meet  $x \wedge x'$  and join  $x \vee x'$  for all each pair of elements of  $A$  are defined with  $\succsim$ . For any subset  $A$  of a poset  $X$ , we denote the supremum and infimum of  $A$  by  $\sup A$  and  $\inf A$ , respectively. That is,  $\sup A$  is the least element in  $X$  such that  $\sup A \succsim a$ , for all  $a \in A$ . Similarly,  $\inf A$  is the greatest element in  $X$  such that  $a \succsim \inf A$ , for all  $a \in A$ . A lattice  $X$  is complete if both  $\inf A$  and  $\sup A$  are in  $X$  for any  $A \subseteq X$ . A chain is a totally ordered poset. A poset  $X$  is (countably) lower chain complete if any (countable) chain  $A \subseteq X$  has its infimum in  $X$ . The poset is (countably) upper chain complete if any such chain has its supremum in  $X$ . The poset is (countably) chain complete if it is both upper and lower (countably) chain complete. Let  $X$  and  $Y$  be subsets of  $\mathbb{R}$ . Set  $Y$  dominates  $X$  in the strong set order if for any  $x$  in  $X$  and  $y$  in  $Y$ , we have  $\max\{x, y\}$  in  $Y$  and  $\min\{x, y\}$  in  $X$ .

Let  $\mathcal{M}(X)$  denote the space of probability measures defined on a compact subset of  $X \subset \mathbb{R}^n$ . Although even if  $X$  is a lattice, the poset  $(X, \succsim_{st})$  is not a lattice as pointed in Kamae, Krengel, and O'Brien (1977). For e.g., for  $(\mathbb{R}^2, \succsim_{st})$  let  $p_1 = 0.5(\epsilon_a + \epsilon_b), p_2 = 0.5(\epsilon_a + \epsilon_c), p_3 = 0.5(\epsilon_c + \epsilon_b), p_4 = 0.5(\epsilon_a + \epsilon_d)$ , where  $a = (0, 0), b = (0, 1), c = (1, 0), d = (1, 1)$ . Then, both  $p_3$  and  $p_4$  are supremum, which is a contradiction. However,  $(\mathbb{R}^2, \succsim_{st})$  is chain-complete.<sup>49</sup> A correspondence  $T : X \rightarrow 2^Y$  is upper-hemicontinuous at a point  $x_0 \in X$  if for any sequence  $\{x_n\}_{n \in \mathbb{N}}$  such that  $\{x_n\} \rightarrow x_0, y_n \in T(x_n), y_0 \in T(x_0)$  implies  $y_n \rightarrow y_0$ . We will require the following existence theorem in Smithson (1971) for chain-complete (non-lattice) spaces, and refer the interested reader to his paper for further details.

**Theorem** (Smithson's fixed point theorem (1971)). *Let  $X$  be a chain-complete poset equipped with partial order  $\succsim$ , and  $T : X \rightarrow 2^X$  a Type I (Type II) monotone correspondence. Suppose for any chain  $C$  in  $X$  and any monotone selection  $f$  from the restriction of  $T$  to  $C$ ,  $f : C \rightarrow X$ , there exists  $y_0 \in T(\sup C)$  such that  $y_0 \succsim f(x)$  for all  $x \in C$ . Then if there exists a point  $e \in X$  and a point  $y \in T(e)$  such that  $y \succsim e$ , then  $T$  has a fixed point.*

<sup>49</sup>This also holds true for increasing and convex order,  $\succsim_{icx}$ .

## B Proofs

This section is divided into four parts. Part (i) proves two auxiliary lemmas that are instrumental for our main results. Part (ii) proves Theorems 1 and 2 of the paper and part (iii) proves Theorems 3 and 4, respectively. Part (iv) does the welfare comparison.

The proofs for Theorems 1 and 2 are structured in three steps, with the proof technique similar to that in Acemoglu and Jensen (2015). For Theorem 2, the *first* step involves showing that for any fixed model distribution  $\mu \in \mathcal{M}_1(\Theta)$ , the set of stationary distributions on states and actions,  $m^*$ , induced by the optimal policy correspondence  $G$ , will be Type I (Type II) monotonic in the primitives  $p$ . The *second* step involves constructing a mapping  $\hat{\theta}$  that for each given  $\mu$  and  $p$  yields a set of model distributions,  $\mu'$ s. It is the fixed points of this map that are the equilibrium model distributions  $\mu^*$ , given  $p$ . Finally, the *third* step involves the least and greatest selections from this map will be increasing in  $p$ . This in turn leads us to give a new existence proof of Theorem 1 that relies on the monotonicity and identification properties of the equilibrium map,  $T$ . The rest of the proofs (Theorems 3 and 4) shall follow an analogous structure. Proof of Theorem 5 is based on the Taylor expansion of the utility function and on an application of a new result proven by Canonne (2022) on entropy bounds.

The equilibrium mapping  $T$  associated with the Berk-Nash equilibrium (Definition 5) is a set-valued function on the product space of probability measure on states and actions, and parameter space,  $T : W \rightarrow 2^W$ , where  $W = \mathcal{M}_1(\mathbb{S} \times \mathbb{X}) \times \mathcal{M}_1(\Theta)$  such that  $\mathcal{T}(m, \mu) = \mathcal{M}(m, \mu) \times \mathcal{M}(\Theta_Q(m))$ , where

$$(m, \mu) \mapsto \mathcal{M}(m, \mu) \equiv \left\{ m' \in \mathcal{M}_1(\mathbb{S} \times \mathbb{X}) : m' \in \mathcal{F}(\mu) \ \& \ m'_S(\cdot) = \int_{\mathbb{S} \times \mathbb{X}} Q(\cdot | s, x) m(ds, dx) \right\} \quad (16)$$

for any  $\mu \in \mathcal{M}_1(\Theta)$ ,  $\mathcal{F}(\mu)$  is the set of all  $m'$  that satisfies the condition of optimality and stationarity as in Definition 5.

### (i) Auxiliary Lemmas

Towards this end, we begin by proving two lemmas. They adapt lemmas in Acemoglu and Jensen (2015) that are established for exogenous shock processes to our setting of endogenous Markov decision process. Lemma 1 shows that under assumptions 1 and 2, the optimal policy correspondence, defined in Eq (8),  $G : \mathbb{S} \times P \rightarrow 2^X$  is increasing in the strong-set order and the least and the greatest selection are increasing in  $s$ .

**Lemma 1.** *Let the regular SMDP satisfy assumptions 1 and 2. Then for any given  $p$ , the optimal policy correspondence  $G : \mathbb{S} \times P \rightarrow 2^X$  is increasing in the state  $s$  in the strong set order. In*



particular, it has a least and a greatest selection, and they are increasing in the state  $s$ .<sup>50</sup>

*Proof.* Given a model distribution  $\mu$ , and corresponding to a Bellman solution  $V$  of Eq. (5),

$$G(s, \mu, p) = \operatorname{argmax}_{x \in \mathbb{X}} u(s, x) + \delta \int_S V(s', \mu, p) \bar{Q}_\mu(ds'|s, x), \quad (17)$$

where  $\bar{Q}_\mu = \int_{\Theta} Q_\theta \mu(d\theta)$ . Given Assumptions 1 and 2, and by the application of Theorem 3.9.2 in Topkis (1998), the expression on the right-hand side is supermodular in  $(s, x)$ . Therefore, the optimal policy correspondence will be increasing in  $s$  in the strong-set order. Further, given the optimal correspondence, the greatest and least selection exist and are increasing in the state. ■

The next lemma transfers the monotonic nature of the optimal policy correspondence to its corresponding fixed point map,  $T$ , in terms of Type I (Type II) monotonicity.

**Lemma 2.** *If the optimal policy correspondence  $G : \mathbb{S} \rightarrow 2^{\mathbb{X}}$  has an increasing greatest (least) selection, then the fixed point correspondence  $T$  is Type I (Type II) monotone with respect to  $\succeq_{st}$ . Further, if  $G$  depends on a primitive  $p \in P$  such that  $G : \mathbb{S} \times P \rightarrow 2^{\mathbb{X}}$  and the greatest (least) selection from  $G$  is increasing in  $p$ , then the fixed point correspondence  $T$  indexed by  $p$ ,  $T_p$  is Type I (Type II) monotone in  $\mathcal{M}_1(\mathbb{S} \times \mathbb{X})$  with respect to  $\succeq_{st}$ .*

*Proof.* We only prove the above statement for the greatest selection in the Type I case.<sup>51</sup> Consider probability measures,  $\nu_1, \nu_2 \in \mathcal{M}(\mathbb{S} \times \mathbb{X})$  such that  $\nu_2 \succeq_{st} \nu_1$ . To prove that the fixed point correspondence  $T$  is Type I monotone, we need to show that for any  $\lambda_1 \in T\nu_1$ , there exists a  $\lambda_2 \in T\nu_2$  such that  $\lambda_2 \succeq_{st} \lambda_1$ . That is, if  $\lambda_1 \in T\nu_1$ , then there exists a measurable selection  $g_1 : \mathbb{S} \rightarrow \mathbb{X}$ , such that, for all  $A \times B \in \mathcal{B}(\mathbb{S} \times \mathbb{X})$ ,

$$\lambda_1(A, B) = \int_{\mathbb{S} \times \mathbb{X}} Q(A|s, g_1(s)) \chi_B(g_1(s)) \nu_1(ds, dx)$$

and therefore, there must exist a selection  $g_2$  such that  $\lambda_2 \in T\nu_2$ , and  $\lambda_2 \succeq_{st} \lambda_1$ . From Lemma 1, we know that such a greatest selection exists, and therefore, for all  $A \times B \in \mathcal{B}(\mathbb{S} \times \mathbb{X})$ ,

$$\begin{aligned} \lambda_2 \succeq_{st} \lambda_1 &\iff \int_{\mathbb{S}} f(s, g_2(s)) Q(ds|s, g_2(s)) \nu_2(s) \geq \int_{\mathbb{S}} f(s, g_2(s)) Q(ds|s, g_2(s)) \nu_1(s) \\ &\geq \int_{\mathbb{S}} f(s, g_1(s)) Q(ds|s, g_1(s)) \nu_1(s) \end{aligned}$$

The first inequality follows from  $\nu_2 \succeq_{st} \nu_1$ . The second inequality follows from  $Q$  being monotone and  $g_2$  being the greatest selection. With an identical argument, one can prove for

<sup>50</sup>We suppress the dependence on  $\mathcal{M}_1(\Theta)$  for notational convenience.

<sup>51</sup>The rest of the cases follow analogously.

the case where the greatest selection from  $G$  is to be increasing in  $p$ , and therefore,  $T$  indexed by  $p$ ,  $T_p$  is Type I (Type II) monotone in  $\mathcal{M}_1(\mathbb{S} \times \mathbb{X})$ , with respect to  $\succeq_{st}$ . ■

We are now ready to prove Theorems 1 and 2. Throughout the proof, primitive  $p$  is restricted to the set  $P \equiv \{p_1, p_2\}$  ordered by  $p_2 \succeq p_1$ .

(ii) **Proofs for Theorems 1 and 2:** The *first* step involves showing for a given model distribution  $\mu \in \mathcal{M}_1(\Theta)$ , the set of Berk-Nash equilibrium  $m$  induced by the optimal policy correspondence  $G$  will be Type I (Type II) increasing in the primitives  $p \in P$ . To this end, we first show that under a positive shock, the fixed point correspondence  $T$  is Type I (Type II) increasing in the primitives  $p$ . The fixed points of this correspondence are the Berk-Nash equilibrium  $m$  for a given model distribution  $\mu$ , and finally using a result in Acemoglu and Jensen (2015), we show that the set of Berk-Nash equilibrium for a given model distribution is Type I (Type II) increasing in the primitives.

Because of Lemma 1, the stationary optimal policy correspondence  $G$  will have a least and a greatest selection that will be increasing in  $s$  and therefore, by Lemma 2, for a given model distribution  $\mu$ , and primitive  $p$ ,  $T_{\mu,p} : \mathcal{M}_1(\mathbb{S} \times \mathbb{X}) \rightarrow 2^{\mathcal{M}_1(\mathbb{S} \times \mathbb{X})}$  induced by the stationary optimal policy correspondence,  $G$ , defined in (8) is Type I (Type II) monotone with respect to  $\succeq_{st}$ . Now from a routine modification of Theorem B3 in Acemoglu and Jensen (2015),<sup>52</sup> the set of fixed points,  $F : \mu \times P \rightarrow 2^{\mathcal{M}_1(\mathbb{S} \times \mathbb{X})}$ , given by  $F(\mu, p) = \{m \in \mathcal{M}_1(\mathbb{S} \times \mathbb{X}) : m \in T_{\mu,p}m\}$  is non-empty valued and upper hemicontinuous. The proof of upper-hemicontinuity of  $T_{\mu,p}$  follows for the finite case follows from Claim B (Page 744) in EP, while for the infinite case, it follows by invoking the conditions in Definition 3 on a regular SMDP, as proved in Anderson, Duanmu, Ghosh, and Khan (2023). From Lemma 2,  $T_{\mu,p}$  is Type I (Type II) monotone in  $p$ , and therefore, from the monotonicity theorem of Acemoglu and Jensen (2015) in the main text, the set of fixed points  $F$  will be non-empty and Type I (Type II) monotone in  $p$ . Hence, the set of Berk-Nash equilibrium  $m$  induced by the optimal policy correspondence  $G$  will be Type I (Type II) increasing in the primitives  $p$ . This completes the first step.

The *second* step constructs a mapping  $\hat{\theta}$  that for each given model distribution  $\mu$  and primitive  $p \in P$ , yields a set of best-fit model distributions,

$$\hat{\theta}(\mu, p) = \{\theta(m) \equiv \operatorname{argmin}_{\theta \in \Theta} K_Q(m, \theta) : m \in F(\mu, p)\}. \quad (18)$$

A model distribution  $\mu^*$  is an equilibrium belief if and only if  $\mu^* \in \hat{\theta}(\mu^*, p)$ . By assumption

---

<sup>52</sup>Theorem B3 in Acemoglu and Jensen (2015) builds on the existence theorem in Smithson (1971). It states the following: Assume that the equilibrium mapping  $T$  is either Type I or type II monotone. In addition, assume that the set of measures on state and actions has an infimum. Then  $T$  has fixed point. In addition, the fixed-point correspondence is upper hemicontinuous if  $T$  is upper hemicontinuous.

3 of point identification and therefore, by uniqueness, for each  $m$ , there will be a non-empty unique  $\theta(m)$ . That is, one has a Dirac measure on  $\theta(m)$ . From Berge's Maximum Theorem,  $\theta$  is continuous<sup>53</sup> and given the upper hemicontinuity of  $F(\mu, p)$ ,  $\hat{\theta}$  will be upper hemi-continuous. The rest of this step follows Acemoglu and Jensen (2015). For a fixed  $\mu$ , and given  $F(\mu, \cdot)$  is Type I and Type II monotone in  $p$ , and increasing  $\theta(m)$  under Assumption 4, one can use their Theorem 4 (pp. 601) to conclude that the least and greatest selections from  $\hat{\theta}(\mu, \cdot)$  will be increasing in  $p$  holding  $\mu$  fixed.<sup>54</sup>

The *third* and *final* step shows that if the least and the greatest selections of the upper-hemicontinuous fixed point correspondence  $\hat{\theta}$  are increasing in  $p$ , then the fixed points are increasing in  $p$ . Since under Assumption 4,  $\theta$  is a monotonic function of  $m$ , therefore,  $\mu_{\min} \equiv \theta(\delta_{\inf \mathbb{S} \times \mathbb{X}})$  and  $\mu_{\max} \equiv \theta(\delta_{\sup \mathbb{S} \times \mathbb{X}})$ , where  $\delta_{\mathbb{S} \times \mathbb{X}}$  denotes the degenerate measure on  $\mathbb{S} \times \mathbb{X}$  with its mass at  $(s, x)$ . Therefore,  $\mu \succeq \mu_{\min}$  for all  $\mu \in \hat{\theta}(\mu_{\min})$  and  $\mu \preceq \mu_{\max}$  for all  $\mu \in \hat{\theta}(\mu_{\max})$ . Therefore, for every  $p \in P$ ,  $\hat{\theta}(\cdot, p) : [\mu_{\min}, \mu_{\max}] \rightarrow 2^{[\mu_{\min}, \mu_{\max}]}$ . Notice in step 2 that  $F(\mu, p)$  is a convex-valued set of fixed points since  $T_{\mu, p}$  is convex-valued; the proof of convex-valuedness follows from EP. Therefore, the set of fixed points from  $F$  is convex-valued and given  $\theta(m)$  is a continuous function,  $\hat{\theta}(\mu, p)$  is convex-valued. Hence, from Acemoglu and Jensen (2013),  $\hat{\theta}(\cdot, p) : [\mu_{\min}, \mu_{\max}] \rightarrow 2^{[\mu_{\min}, \mu_{\max}]}$  is upper hemicontinuous and convex valued and for each fixed value of  $\mu \in [\mu_{\min}, \mu_{\max}]$ , has least and greatest selections and are increasing in  $p$  and from Corollary 2 in Milgrom and Roberts (1994) the least and greatest fixed points  $\mu^* \in \hat{\theta}(\mu^*, p)$  will be increasing in  $p$ . Hence, we have proven that the least and greatest inferred models are increasing in the primitives under a positive shock to a regular SMDP. From Theorem 2.8.3 in Topkis (1998), and by treating a change (positive/negative) in  $\mu$  as a change in primitive, the remaining part of Theorem 2 follows.

Existence is yielded in Step 2 by the Kakutani-Fan-Glicksberg Theorem since our map  $\hat{\theta}(\mu, p)$  is convex-valued, upper hemi-continuous, and as is shown in EP,  $\mathcal{M}_1(\Theta)$  are locally convex Hausdorff spaces. ■

(iii) **Proof for Theorem 3:** The parameter space  $\Theta \subseteq \mathbb{R}$  is one-dimensional, and hence, the weighted KL divergence satisfies quasimodularity by triviality on the parameter space,  $\Theta \subseteq \mathbb{R}$ . Hence, from Topkis (1978) and Milgrom and Shannon (1994), given an increase in the parameter space from  $\Theta_1$  to  $\Theta_2$  in the strong-set order, the set of minimizers are increasing in the strong-set

<sup>53</sup>The function  $\theta_Q(m)$  is continuous if it is continuous in the weak\* topology on its domain.

<sup>54</sup>Theorem 4 in Acemoglu and Jensen (2015) guarantees least and greatest selections from a fixed point correspondence.

order, i.e.,

$$\Theta(m; \Theta_1) \equiv \operatorname{argmin}_{\theta \in \Theta_1} K_Q(m, \theta) \subseteq \operatorname{argmin}_{\theta \in \Theta_2} K_Q(m, \theta) \equiv \Theta(m; \Theta_2)$$

Therefore, if a change in model distribution is a positive shock then the proof follows on the lines of Theorem 2. Thus, this completes the proof.  $\blacksquare$

**Proof for Theorem 4:** Consider probability measures,  $\nu_1, \nu_2 \in \mathcal{M}(\mathbb{S} \times \mathbb{X})$  such that  $\nu_2 \succeq_{icx} \nu_1$ . The proof of Type I and Type II monotonicity is similar to that of Theorems 1 and 2, albeit requires that we work with increasing convex orders on the set of states and actions.

To prove that the fixed point correspondence  $T$  is Type I monotone in  $\succeq_{icx}$ , we need to show that for any  $\lambda_1 \in T\nu_1$ , there exists a  $\lambda_2 \in T\nu_2$  such that  $\lambda_2 \succeq_{icx} \lambda_1$ , that is, if  $\lambda_1 \in T\nu_1$ , then there exists a measurable selection  $g_1 : \mathbb{S} \rightarrow \mathbb{X}$ , such that, for all  $A \times B \in \mathcal{B}(\mathbb{S} \times \mathbb{X})$ ,

$$\lambda_1(A, B) = \int_{\mathbb{S} \times \mathbb{X}} Q(A|s, g_1(s)) \chi_B(g_1(s)) \nu_1(ds, dx)$$

and therefore, there must exist a selection  $g_2$  such that  $\lambda_2 \in T\nu_2$ , and  $\lambda_2 \succeq_{icx} \lambda_1$ . From Lemma 1, we know that such a greatest selection exists, and therefore, for all  $A \times B \in \mathcal{B}(\mathbb{S} \times \mathbb{X})$ ,

$$\begin{aligned} \lambda_2 \succeq_{icx} \lambda_1 &\iff \int_{\mathbb{S}} f(s, g_2(s)) Q(ds|s, g_2(s)) \nu_2(s) \geq \int_{\mathbb{S}} f(s, g_2(s)) Q(ds|s, g_2(s)) \nu_1(s) \\ &\geq \int_{\mathbb{S}} f(s, g_1(s)) Q(ds|s, g_1(s)) \nu_1(s) \end{aligned}$$

The first inequality follows from  $\nu_2 \succeq_{icx} \nu_1$ . The second inequality follows from  $Q$  being monotone and  $g_2$  being the greatest selection. The only adjustment in the proof that remains is that the best-fit set construction needs to incorporate Assumption 5 for it to have them monotonic, i.e., the *second* step constructs a mapping  $\hat{\theta}$  that for each given model distribution  $\mu$  and primitive  $p \in P$ , yields a set of model distributions,

$$\hat{\theta}(\mu, p) = \{\theta(m) \equiv \operatorname{argmin}_{\theta \in \Theta} K_Q(m, \theta) : m \in F(\mu, p)\}. \quad (19)$$

Hence, this completes the proof.  $\blacksquare$

(iv) **Proof for Theorem 5:** The first part follows from the optimality of  $g(s, \theta^*)$  over  $g(s, \theta_*)$ , under the objective welfare function  $W(s, \theta^*)$ , where the expectation is taken with respect to the true transition function,  $Q$ . Next, under correct learning parameterized by  $\theta^*$ , and for an

initial state  $s_0$ , the objective welfare is given as,

$$W(s_0, \theta^*) = \mathbb{E}_{Q(\cdot|s_{t-1}, g(s_{t-1}, \theta^*))} \left[ \sum_{t=0}^{\infty} \beta^t u(s_t, g(s_t, \theta^*)) \right], \quad t = 0, 1, 2, \dots \quad (20)$$

Unwrapping the above expression,

$$\begin{aligned} W(s_0, \theta^*) &= u(s_0, g(s_0, \theta^*)) + \beta \int_S u(s_1, g(s_1, \theta^*)) Q(ds_1|s_0, g(s_0, \theta^*)) + \dots \\ &\quad \beta^t \int_S u(s_t, g(s_t, \theta^*)) Q(ds_t|s_{t-1}, g(s_{t-1}, \theta^*)) + \dots \end{aligned} \quad (21)$$

Now, with the Berk-Nash parameter  $\theta_*$ , the corresponding objective welfare under misspecified learning is,

$$\begin{aligned} W(s_0, \theta_*) &= u(s_0, g(s_0, \theta_*)) + \beta \int_S u(s_1, g(s_1, \theta_*)) Q(ds_1|s_0, g(s_0, \theta_*)) + \dots \\ &\quad \beta^t \int_S u(s_t, g(s_t, \theta_*)) Q(ds_t|s_{t-1}, g(s_{t-1}, \theta_*)) + \dots \end{aligned} \quad (22)$$

Our objective is to find an upper bound for  $\|W(s_0, \theta^*) - W(s_0, \theta_*)\|$ , where  $\|\cdot\|$  denotes the sup-norm in the function space. Comparing the first terms in Equations (21) and (22) and by a Taylor expansion of  $u(s_0, g(s_0, \theta^*))$  on  $u(s_0, g(s_0, \theta_*))$ , given that we assume that it is continuously differentiable and strictly concave, we have,

$$\|u(s_0, g(s_0, \theta^*)) - u(s_0, g(s_0, \theta_*))\| \leq \left| \frac{du}{dg} \Big|_{g(s_0, \theta_*)} \right| \|g(s_0, \theta^*) - g(s_0, \theta_*)\|.$$

The inequality follows from the concavity of  $u$  in policy function  $g$  and from taking the sup-norm on both sides. Similarly, for the second term, we have,

$$\begin{aligned} &\beta \left\| \left( \int_S u(s_1, g(s_1, \theta^*)) Q(ds_1|s_0, g(s_0, \theta^*)) - \int_S u(s_1, g(s_1, \theta_*)) Q(ds_1|s_0, g(s_0, \theta_*)) \right) \right\| \\ &\leq \beta \left( \int_S |u(s_1, g(s_1, \theta_*))| \cdot \|Q(ds_1|s_0, g(s_0, \theta^*)) - Q(ds_1|s_0, g(s_0, \theta_*))\| + \left| \frac{du}{dg} \Big|_{g(s_1, \theta_*)} \right| \|g(s_1, \theta^*) - g(s_1, \theta_*)\| \right) \end{aligned} \quad (23)$$

This again follows from the strict concavity of  $u$ , a corresponding Taylor-expansion of the utility function  $u$ , and the fact that the integral of the density function (always positive) is 1.

While we have corresponding upper bounds for the utility function, its first derivative, and the difference in the policy functions under the two learning instances, we still need to establish a corresponding upper bound for the total-variation (TV) norm for densities  $Q$ . Towards this end, we rely on a lemma by Canonne (2022), that provides an almost stricter bound on the

TV-norm between densities.<sup>55</sup>

**Theorem** (Bretagnolle–Huber bound, Canonne (2022)). *For any two probability distribution functions  $Q_1, Q_2$  over  $\mathbb{R}$ ,*

$$d_{TV}(Q_1, Q_2) = \|Q_1 - Q_2\|_1 \leq 2\sqrt{1 - e^{-KL(Q_1\|Q_2)}},$$

where  $d_{TV}$  is the total variation norm between densities  $Q_1$  and  $Q_2$ , and  $KL$  is the relative entropy between  $Q_1$  and  $Q_2$ .

Therefore, the first part of the right-hand side expression in Eq.(23) can be bounded as,

$$\beta \left( \int_S |u(s_1, g(s_1, \theta_*))| \cdot \|Q(ds_1|s_0, g(s_0, \theta^*)) - Q(ds_1|s_0, g(s_0, \theta_*))\| \right) \leq 2\beta m_0 \sqrt{1 - e^{-KL(\gamma, Q)}}.$$

where the KL distance,  $KL_1(\gamma, Q)$ , depends on the approximation error  $\gamma$ , and the true transition function,  $Q$ , and also the particular state realization. For the second part, we have,

$$\left| \frac{du}{dg} \Big|_{g(s_1, \theta_*)} \|g(s_1, \theta^*) - g(s_1, \theta_*)\| \right| \leq m_1 \gamma. \quad (24)$$

By repeating the above step for all the terms in  $W(s, \theta^*)$  and  $W(s, \theta_*)$ , we have,

$$\begin{aligned} \|W(s, \theta^*) - W(s, \theta_*)\| &\leq m_1 \gamma + \beta(m_1 \gamma + 2m_0 \sqrt{1 - e^{-KL_1(\gamma, Q)}}) + \beta^2(m_1 \gamma + 2m_0 \sqrt{1 - e^{-KL_2(\gamma, Q)}}) + \dots \\ &\leq \frac{m_1 \gamma + 2m_0 \beta \sqrt{1 - e^{-KL_*(\gamma, Q)}}}{1 - \beta} \end{aligned}$$

where  $KL_*(\gamma, Q)$  is the upper bound on the sequence of KL distances,  $\{KL_1, KL_2, \dots, KL_n, \dots\}$ . Hence, we have established our upper bound between the two welfare quantities, and this completes the proof. ■

---

<sup>55</sup>Our choice of this particular bound over the well-known Pinsker bound is that as the KL divergence gets larger, the Pinsker bound becomes vacuous and exceeds the trivial bound of 1. The BH bound, however, never exceeds 1.

## References

- ACEMOGLU, D., AND M. K. JENSEN (2013): “Aggregate comparative statics,” *Games and Economic Behavior*, 81, 27–49.
- (2015): “Robust comparative statics in large dynamic economies,” *Journal of Political Economy*, 123(3), 587–640.
- AMIR, R. (1996): “Sensitivity analysis of multisector optimal economic dynamics,” *Journal of mathematical economics*, 25(1), 123–141.
- (2018): “Supermodularity and monotone methods in economics,” .
- ANDERSON, R. M., H. DUANMU, A. GHOSH, AND M. A. KHAN (2023): “On existence of Berk-Nash equilibria in misspecified Markov decision processes with infinite spaces,” *arXiv preprint arXiv:2206.08437*.
- ARROW, K. J., D. BLACKWELL, AND M. A. GIRSHICK (1949): “Bayes and minimax solutions of sequential decision problems,” *Econometrica*, pp. 213–244.
- ARROW, K. J., AND J. R. GREEN (1973): “Notes on expectations equilibria in Bayesian settings,” *Institute for Mathematical Studies in the Social Sciences*.
- ARROW, K. J., T. HARRIS, AND J. MARSCHAK (1951): “Optimal inventory policy,” *Econometrica*, pp. 250–272.
- BALBUS, Ł., P. DZIEWULSKI, K. REFFETT, AND Ł. WOŹNY (2022): “Markov distributional equilibrium dynamics in games with complementarities and no aggregate risk,” *Theoretical Economics*, 17(2), 725–762.
- BERK, R. H. (1966): “Limiting behavior of posterior distributions when the model is incorrect,” *The Annals of Mathematical Statistics*, 37(1), 51–58.
- BOHREN, J. A., AND D. N. HAUSER (2021): “Learning with heterogeneous misspecified models: Characterization and robustness,” *Econometrica*, 89(6), 3025–3077.
- BROCK, W. A., AND L. J. MIRMAN (1972): “Optimal economic growth and uncertainty: the discounted case,” *Journal of Economic Theory*, 4(3), 479–513.
- BUNKE, O., AND X. MILHAUD (1998): “Asymptotic behavior of Bayes estimates under possibly incorrect models,” *The Annals of Statistics*, 26(2), 617–644.
- CANONNE, C. L. (2022): “A short note on an inequality between KL and TV,” *arXiv preprint arXiv:2202.07198*.
- CHE, Y.-K., J. KIM, AND F. KOJIMA (2021): “Weak monotone comparative statics,” *arXiv preprint arXiv:1911.06442*.
- COVER, T. M., AND J. THOMAS (2005): *Elements of information theory*. John Wiley & Sons.
- CSABA, D., AND B. SZOKE (2023): “Learning with misspecified models,” Discussion paper, mimeo.
- DATTA, M., K. REFFETT, AND Ł. WOŹNY (2018): “Comparing recursive equilibrium in

- economies with dynamic complementarities and indeterminacy,” *Economic Theory*, 66, 593–626.
- DEIMEN, I., AND J. WIRTZ (2022): “Control, cost, and confidence: Perseverance and procrastination in the face of failure,” *Games and Economic Behavior*, 134, 52–74.
- DEN HAAN, W. J., AND T. DRECHSEL (2019): “Misspecification in macroeconomics: Difficult but not impossible to deal with,” .
- ESPONDA, I., AND D. POUZO (2015): “Equilibrium in misspecified Markov decision processes,” *arXiv preprint arXiv:1502.06901*.
- (2016): “Berk–Nash equilibrium: A framework for modeling agents with misspecified models,” *Econometrica*, 84(3), 1093–1130.
- (2021): “Equilibrium in misspecified Markov decision processes,” *Theoretical Economics*, 16, 717–757.
- ESPONDA, I., D. POUZO, AND Y. YAMAMOTO (2021): “Asymptotic behavior of Bayesian learners with misspecified models,” *Journal of Economic Theory*, 195, 105260.
- FARMER, L., E. NAKAMURA, AND J. STEINSSON (2023): “Learning about the long run,” Discussion paper, NBER Working Paper No. 29495.
- FRICK, M., R. IJIMA, AND Y. ISHII (2020): “Misinterpreting others and the fragility of social learning,” *Econometrica*, 88(6), 2281–2328.
- GRÜNWARD, P., AND T. VAN OMMEN (2017): “Inconsistency of Bayesian Inference for Misspecified Linear Models, and a Proposal for Repairing It,” *Bayesian Analysis*, 12(4), 1069 – 1103.
- HANSEN, L. P., AND T. J. SARGENT (2011): *Robustness*. Princeton university press.
- HEIDHUES, P., B. KÓSZEGI, AND P. STRACK (2018): “Unrealistic expectations and misguided learning,” *Econometrica*, 86(4), 1159–1214.
- HÖFT, H. (1987): “Order preserving selections for multifunctions,” *Contributions to General Algebra*, 5.
- HOPENHAYN, H. A., AND E. C. PRESCOTT (1992): “Stochastic monotonicity and stationary distributions for dynamic economies,” *Econometrica*, pp. 1387–1406.
- HUBER, P. J. (1967): “The behavior of maximum likelihood estimates under nonstandard conditions,” in *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, vol. 1, pp. 221–233. Berkeley, CA: University of California Press.
- HUGGETT, M. (2003): “When are comparative dynamics monotone?,” *Review of Economic Dynamics*, 6(1), 1–11.
- JEHIEL, P. (2005): “Analogy-based expectation equilibrium,” *Journal of Economic theory*, 123(2), 81–104.
- KAMAE, T., U. KRENGEL, AND G. L. O’BRIEN (1977): “Stochastic inequalities on partially ordered spaces,” *The Annals of Probability*, 5(6), 899–912.



- KIRMAN, A. P. (1975): “Learning by firms about demand conditions,” in *Adaptive economic models*, pp. 137–156. Elsevier.
- KOULOVIATIANOS, C., L. J. MIRMAN, AND M. SANTUGINI (2009): “Optimal growth and uncertainty: Learning,” *Journal of Economic Theory*, 144(1), 280–295.
- KYDLAND, F. E., AND E. C. PRESCOTT (1982): “Time to build and aggregate fluctuations,” *Econometrica*, pp. 1345–1370.
- LANZANI, G. (2022): “Dynamic concern for misspecification,” *mimeo*.
- LEVHARI, J. D., AND T. N. SRINIVASAN (1969): “Optimal savings under uncertainty,” *The Review of Economic Studies*, 36(2), 153–163.
- LEWBEL, A. (2019): “The identification zoo: Meanings of identification in econometrics,” *Journal of Economic Literature*, 57(4), 835–903.
- LIGHT, B. (2021): “Stochastic Comparative Statics in Markov Decision Processes,” *Mathematics of Operations Research*.
- LONG JR., J. B., AND C. I. PLOSSER (1983): “Real business cycles,” *Journal of political Economy*, 91(1), 39–69.
- LUCAS, R. E. (1978): “Asset prices in an exchange economy,” *Econometrica*, pp. 1429–1445.
- LUCAS, R. E., AND E. C. PRESCOTT (1971): “Investment under uncertainty,” *Econometrica*, pp. 659–681.
- (1974): “Equilibrium search and unemployment,” *Journal of Economic Theory*, 7(2), 188–209.
- MILGROM, P., AND J. ROBERTS (1994): “Comparing equilibria,” *The American Economic Review*, pp. 441–459.
- MILGROM, P., AND C. SHANNON (1994): “Monotone comparative statics,” *Econometrica*, pp. 157–180.
- MOLAVI, P. (2019): “Macroeconomics with learning and misspecification: A general theory and applications,” *Unpublished manuscript*.
- NYARKO, Y. (1991): “Learning in mis-specified models and the possibility of cycles,” *Journal of Economic Theory*, 55(2), 416–427.
- PHELPS, E. S. (1962): “The new view of investment: a neoclassical analysis,” *The Quarterly Journal of Economics*, 76(4), 548–567.
- PUTERMAN, M. L. (1994): *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.
- ROSS, W. (1987): “The expectation of the likelihood ratio criterion,” *International Statistical Review/Revue Internationale de Statistique*, pp. 315–330.
- RUST, J. (1994): “Structural estimation of Markov decision processes,” *Handbook of econometrics*, 4, 3081–3143.

- SAGHAFIAN, S. (2018): “Ambiguous partially observable Markov decision processes: Structural results and applications,” *Journal of Economic Theory*, 178, 1–35.
- SANTOS, M. S. (2000): “Accuracy of numerical solutions using the Euler equation residuals,” *Econometrica*, 68(6), 1377–1402.
- SHAKED, M., AND J. G. SHANTHIKUMAR (2007): *Stochastic orders*. Springer.
- SHALIZI, C. R. (2009): “Dynamics of Bayesian updating with dependent data and misspecified models,” *Electronic Journal of Statistics*, 3, 1039–1074.
- SMITHSON, R. (1971): “Fixed points of order preserving multifunctions,” *Proceedings of the American Mathematical Society*, 28(1), 304–310.
- TOPKIS, D. M. (1978): “Minimizing a submodular function on a lattice,” *Operations research*, 26(2), 305–321.
- (1998): *Supermodularity and complementarity*. Princeton university press.
- TORRES, R. (2005): “Multivariate Monotone Bayesian Updating,” *Available at SSRN 2167744*.
- UPPAL, R., AND T. WANG (2003): “Model misspecification and underdiversification,” *The Journal of Finance*, 58(6), 2465–2486.
- WHITE, H. (1982): “Maximum likelihood estimation of misspecified models,” *Econometrica*, pp. 1–25.