The Ronald O. Perelman Center for Political Science and Economics (PCPSE)
133 South 36th Street
Philadelphia, PA 19104-6297

pier@econ.upenn.edu
http://economics.sas.upenn.edu/pier

PIER Working Paper
21-018

Robust Model Misspecification and Paradigm Shifts

CUIMIN BA
University of Pennsylvania

June 23, 2021

https://ssrn.com/abstract=3914106

Robust Model Misspecification and Paradigm Shifts

Cuimin Ba*

First version: Apr 14, 2021

This version: June 23, 2021

– Click here to see the most recent version –

Abstract

This paper studies the forms of model misspecification that are likely to persist when compared with competing models. I consider an agent using a subjective model to learn about an action-dependent outcome distribution. Aware of potential model misspecification, she uses a threshold rule to switch between models according to how well they fit the data. A model is *globally robust* if it can persist against every finite set of competing models and is *locally robust* if it can persist against every finite set of nearby competing models. The main result provides simple characterizations of globally robust and locally robust models based on the set of Berk-Nash equilibria they induce. I then apply the results to examples including risk underestimation, overconfidence, and incorrect beliefs about market demand.

Keywords: misspecified Bayesian learning, competing models, robust misspecification, Berk-Nash equilibrium

*University of Pennsylvania. Email: cuiminba@sas.upenn.edu. I am deeply indebted to Aislinn Bohren and George Mailath for their guidance and support at every stage of this paper. I thank Alvaro Sandroni, Kevin He, Yuhta Ishii, Andrew Postlewaite, Juuso Toikka, Hanming Fang, Changhwa Lee, and several conference and seminar participants for helpful comments and suggestions.

1 Introduction

Economists have long incorporated the idea of subjective models into their modeling of economic agents. The recent literature on misspecified learning explores the behavioral and welfare implications of using incorrect models. Depending on the forms of misspecification, learners may not learn the true state of the world and thus may react suboptimally.[1]

The assumption that individuals forever hold on to a single misspecified model is questionable. There is a plethora of evidence suggesting that individuals look for better alternatives and switch between models. Take economists and data scientists for example. They use a specific econometric model, estimate the parameters, and make policy recommendations accordingly. However, they often switch when an alternative model seems to better fit the data, such as including a set of new explanatory variables in a regression, or accepting the Natural Rate Hypothesis in place of the Phillips Curve. The philosophy of science also offers numerous examples of paradigm shifts in scientific advances (Kuhn, 1962). There is evidence that even non-experts in statistics can have multiple subjective models and switch to another model if necessary. For example, people are influenced by and attracted to different narratives or political views as they receive more information (Fisher, 1985; Braungart and Braungart, 1986). They also strive for overcoming their implicit bias through self-reflection (Wegener and Petty, 1997; Massey and Wu, 2005; Di Stefano, Gino, Pisano, and Staats, 2015).

If decision makers entertain competing models, when should we expect them to keep their current misspecified model? Which forms of misspecification are more likely to persist? In other words, when are subjective models *robust*? This paper proposes a framework of misspecified Bayesian learning that allows agents to revise their models.

I consider an infinite-period decision problem of a single agent. In each period, the agent chooses an action and then observes an outcome, the distribution of which is unknown to the agent and contingent on the action. The agent then obtains a flow payoff jointly determined by the action and the realized outcome. In contrast to a *dogmatic modeler* who relies on a single model, I consider a *switcher* who switches between models. In particular, she starts

[1]Examples include: a monopolist trying to estimate the slope of the demand function when the true slope lies outside of the support of his prior (Nyarko, 1991; Fudenberg, Romanyuk, and Strack, 2017); agents learning from private signals and other individuals' actions while neglecting the correlation between the observed actions (Eyster and Rabin, 2010; Ortoleva and Snowberg, 2015; Bohren, 2016) or overestimating how similar others' preferences are to their own (Gagnon-Bartsch, 2017); overconfident agents falsely attributing low outcomes to an adverse environment (Heidhues, Kőszegi, and Strack, 2018, 2019); a decision maker imposing false causal interpretations on observed correlations (Spiegler, 2016, 2019, 2020); a gambler who flips a fair coin mistakenly believing that future tosses must exhibit systematic reversal (Rabin and Vayanos, 2010; He, 2020); individuals narrowly focusing their attention on only a few aspects rather than a complete state space (Mailath and Samuelson, 2020).

with an *initial* model, while simultaneously considering a set of *competing* models as a potential replacement. The models are parametric: a model, together with a specific parameter value, corresponds to a particular profile of action-dependent outcome distributions. In each period, she either keeps the model she used last period or switches to an alternative. The agent uses the current model to complete two tasks: reasoning and acting. That is, she uses Bayesian inference to update her belief about its parameters and chooses the action that maximizes her discounted sum of payoffs given her posterior derived from this model.

The framework clarifies the distinction between learning within a model (updating beliefs over parameters) and identifying which model to use (model switching or "paradigm shifts"). In order to have a disciplined way to make decisions, our agent, despite having concerns about model misspecification, uses an initial subjective model to interpret the world and guide her actions, until the data reveals to her the superiority of a competing model in describing the world. When she remains under the same paradigm, she behaves exactly as prescribed by the theory of subjective probability (Savage, 1972), evaluating the probability of each (subjective) state of the world using the Bayes rule and behaving in a dynamically consistent manner. Here, each state corresponds to a parameter value while the potentially incomplete state space corresponds to a model. The incompleteness of the state space may stem from the coarseness of human thinking, a constraint on cognitive ability, or simply a lack of information and knowledge. I capture the conceptual distinction between updating models and changing models by the different learning rules: while the agent performs Bayesian learning about the parameters within a model, she switches to a different model only if it is compelling enough according to a *Bayes factor criterion* (Kass and Raftery, 1995). The Bayes factor for a model is the likelihood ratio of outcomes under this model and the model used in the last period. The agent switches to the competing model that generates the highest likelihood ratio if this ratio is above an exogenous switching threshold that is larger than 1 and does not switch if all likelihood ratios are below the threshold.

I develop three different notions of robustness based on the long-term persistence possibility of models. A model is said to *persist* against a set of competing models if, when given as an initial model, the agent eventually adopts it forever with positive probability. A model is *globally robust* if for every set of competing models, there exists a prior and a policy for the model under which this model persists. Global robustness is a good criterion to determine which forms of misspecification are more likely to persist, because this global notion requires a positive chance of persistence no matter what competing models are entertained. However, when the agent has a limited understanding of the world or when the agent is reluctant to consider significant changes, it makes sense to restrict our attention to competing models that are close to the initial model, as the agent may only take small steps in her exploration of new models. In light of this, I propose two notions of local robustness. A model is *unconstrained locally robust* if for every set of "neighbor" competing models, there exists a prior and a policy such that this model persists. Here, models are nearby each other if they predict sufficiently similar outcome distributions. This notion is unconstrained because it places no restriction on the similarity of the parametric structures between models. In contrast, a model is *constrained locally robust* if it can persist against all nearby competing models that belong to the same parametric family. For example, consider a monopolist who is trying to estimate a market demand function that he believes to be linear. An unconstrained locally robust model would be able to persist against all nearby demand functions, potentially complicated and non-linear, as long as they are close to the monopolist's initial conjecture. Constrained local robustness, on the other hand, only requires a model to persist against all nearby linear models, such as a linear demand function with a slightly different intercept.

Several challenges arise when incorporating a model switching process into a misspecified learning problem. Since the agent has access to multiple models, we need to keep track of multiple belief processes. All processes are generated by endogenous data—the agent's action in this period induces a posterior that in turn alters her subsequent play. As is widely recognized in the misspecified learning literature, such belief processes may oscillate forever. Beyond that, the agent's best response changes with the current model choice, further intertwining the learning processes. Consequently, the Bayes factors may fail to have good convergence properties, making it difficult to assess the long-term performance of models.

The main results of this paper characterize each notion of robustness. I first provide a characterization of global robustness (Theorem 1). Notice that a globally robust model must be able to persist against every correctly specified model. This observation resolves the aforementioned difficulty. In fact, when the competing model is correctly specified, the likelihood ratio, as well as all beliefs must converge as an implication of the Martingale Convergence Theorem (Lemma 1). This allows us to compare the goodness-of-fit of models only at the limit beliefs. In a self-confirming equilibrium, the agent holds a belief that exactly matches the objective outcome distribution. A self-confirming equilibrium under a particular model is *p-absorbing* if the action of a dogmatic modeler who only uses this model converges to the support of this self-confirming equilibrium with positive probability. Lemma 2 shows that if a model persists against some correctly specified model, then there must exist a p-absorbing self-confirming equilibrium under this model.

Theorem 1 further establishes that the existence of such an equilibrium is not only necessary but also sufficient for global robustness. This equivalence reduces the complicated problem of a switcher to the problem of a dogmatic modeler. The intuition behind Theorem 1 is as follows: when the agent starts with a prior that is sufficiently close to an equilibrium belief, her model almost perfectly fits the observed data and hence she has no reason to switch, with the p-absorbing condition ensuring that she never deviates from the equilibrium with positive probability. Building on existing results from the literature, Corollary 2 shows that a uniformly quasi-strict self-confirming equilibrium is p-absorbing, thus providing a sufficient condition for global robustness which is straightforward to verify from the primitives.

I next turn to the characterization of unconstrained local robustness. In principle, the set of unconstrained locally robust models can be much larger than the set of globally robust models. However, Theorem 2 reveals that they are actually equivalent. The intuition is simple: given a non-globally-robust model, there is always scope to improve how well it fits the data—such improvements can be local and take the form of a convex combination of the current model and the true data generating process.

Constrained local robustness, on the other hand, is indeed much weaker than robustness and characterized differently. Theorem 3 establishes that a model is constrained locally robust if there exists a p-absorbing Berk-Nash equilibrium that satisfies two additional properties. The first property ensures that when the equilibrium is being played, the model can yield a weakly lower Kullback-Leibler divergence than nearby models from the same parametric family. The second is an identification property that guarantees that a nearby model matching the data equally better—in that they yield the same Kullback-Leibler divergence—must lead to the same belief over the outcome distributions. These properties are similar to but weaker than the requirement of a self-confirming equilibrium in the characterization of global robustness. Provided that the two properties are satisfied, there exists a prior that sufficiently close to the Berk-Nash equilibrium belief such that, on a positive measure of paths, the agent never deviates from the equilibrium and thus no switch will be triggered because the initial model explains the data weakly better than the competing model. Theorems 4 and 5 provide necessary conditions for constrained local robustness. In particular, I show that in two special environments, a model is constrained local robustness only if it gives rise to a Berk-Nash equilibrium that satisfies the first property.

I use these results to study the persistence of misspecification in three applications, all of which lead to suboptimal behavior. When a worker can attribute his underperformance to the outside environment, such as the ability of his coworker, his overconfidence is globally robust and can persist indefinitely. For an investor who systematically underestimates investment risks, while her model is non-robust due to its wrong prediction about the volatility, it's constrained locally robust if she only seeks to match the mean of the return. Finally, I consider a monopolist who has a false prior over the market demand function. His misspecification is not constrained locally robust since there always exists a demand function with a slightly different intercept and slope that yields a strictly better fit at all Berk-Nash equilibria.

The rest of the paper is organized as follows. The next subsection discusses the related literature. Section 2 sets up the model and introduces the model switching framework. Section 3 lays out different variations of Berk-Nash equilibria and self-confirming equilibria that will be useful for the analysis. Section 4 defines and characterizes the three notions of robustness. Section 5 discusses an extension to a stronger notion of global robustness. Section 6 concludes.

Related Literature

This paper builds on the literature of learning with subjective models, including Berk (1966), Easley and Kiefer (1988), Esponda and Pouzo (2016), Bohren and Hauser (2021), Esponda, Pouzo, and Yamamoto (2019), Fudenberg, Lanzani, and Strack (2021) and Frick, Iijima, and Ishii (2020), all of which study asymptotic learning outcomes of dogmatic modelers in relatively general environments. Esponda and Pouzo (2016) first propose the concept of Berk-Nash equilibrium. Further, Esponda et al. (2019) find general conditions for a single agent's action frequency to converge to the Berk-Nash equilibrium using tools from stochastic approximation. Fudenberg et al. (2021) establish that a uniformly strict Berk-Nash equilibrium is uniformly stable in the sense that starting from any prior that is sufficiently concentrated on the Kullback-Leibler minimizers, the dogmatic modeler's action converges to the equilibrium with arbitrarily high probability. This paper contributes to the literature by allowing for model switching and proposing various robustness notions for misspecified models to persist.

This paper is most related to a recent set of papers that explores why certain types of misspecification persist. Olea, Ortoleva, Pai, and Prat (2019) characterize the "winner" model in a contest environment where agents make auction bids based on model-based predictions. With the amount of data being limited, their focus is the trade-off between overfitting and underfitting. Cho and Kasa (2015) also study an agent switching between models but assume a different switching rule. In particular, they assume that the agent always compares her subjective outcome distribution with the empirical realizations. This contrasts the agent in my framework who compares her model to a potentially misspecified alternative. Gagnon-Bartsch, Rabin, and Schwartzstein (2020) study the "attentional stability" of models. They examine a setting where agents realize their model is misspecified if implausible observations emerge but only pay attention to data they deem as relevant given the current model.

Two recent papers approach the problem of which forms of misspecification persist from an evolutionary perspective. Fudenberg and Lanzani (2020) study the evolution dynamics when a small share of a large population mutates to enlarge their subjective models at a Berk-Nash equilibrium. They provide sufficient conditions for a Berk-Nash equilibrium to be robust to invasion. Different from my framework where switching depends on the relative goodness-of-fit of models, they assume that subjective models that induce better performing actions increase their prevalence. He and Libgober (2020) also evaluate competing misspecification based on their expected objective payoffs but examine strategic games where misspecification can lead to beneficial wrong beliefs. Relatedly, Frick, Iijima, and Ishii (2021) also study welfare comparisons of learning biases and find that some biases can outperform Bayesian updating. They focus on a class of learning biases that lead to correct learning and define a bias to be better than another when it leads to higher expected objective payoffs in all decision problems. They characterize this ranking by an efficiency index that quantifies the speed of learning.

A few other papers entertain the similar idea that people have access to multiple models and explore its implications. Mullainathan (2002) presents a model of "categorical thinking" in which people switch between coarse categories and policies discontinuously, resulting in overreaction to news. Ortoleva (2012) proposes and axiomatically characterizes an amendment to Bayes' rule that requires the agent to switch to an alternative upon observing zero-probability events. Karni and Vierø (2013) provide a choice-based decision theory to model a self-correcting agent who can expand his universe of subjective states. Finally, Galperti (2019) and Schwartzstein and Sunderam (2019) extend the idea of alterable subjective models to a persuasion setting and study how a principal could persuade an agent to accept a different worldview.

This paper is also related to the statistics literature in model selection. Statisticians have been interested in the best practices of selecting among models and have developed a number of criteria that differ in their cost of computation and penalty for overfitting, such as Bayes factor, Akaike information criterion (AIC), Bayesian information criterion (BIC), and likelihood-ratio test (LR test). The machine learning community favors cross-validation due to its flexibility and ease of use. All of these criteria are shown to be asymptotically correct under different assumptions (Chernoff, 1954; Akaike, 1974; Stone, 1977; Schwarz et al., 1978; Kass and Raftery, 1995; Konishi and Kitagawa, 2008). This paper focuses on the Bayes factor rule and contributes to the literature by assuming an endogenous data-generating process. I will come back to the comparison of different model selection rules in Section 2.5.

2 Framework

2.1 Objective Environment

A single agent with discount factor $\delta < 1$ makes decisions in an infinitely repeated problem. In each period $t = 0, 1, 2, ...$, the agent chooses an action a_t from a finite action set $\mathcal{A}$ and then observes an outcome y_t from $\mathcal{Y}$, with $\mathcal{Y}$ being either $\mathbb{R}^n$ or a compact subset of $\mathbb{R}^n$ for some positive integer n. Conditional on a_t, outcome y_t is drawn according to probability measure $Q^*(\cdot|a_t) \in \Delta\mathcal{Y}$. This true data generating process (henceforth true DGP) remains fixed throughout. At the end of period t, she obtains a flow payoff $u_t := u(a_t, y_t) \in \mathbb{R}$. Denote the observable history in the beginning of period t by $h_t := (a_\tau, y_\tau)_{\tau=0}^{t-1}$ and the set of all such histories by $H_t = (\mathcal{A} \times \mathcal{Y})^t$.

Assumption 1. *(i) For all $a \in \mathcal{A}$, $Q^*(\cdot|a)$ is absolutely continuous w.r.t. a common measure ν, and the Radon-Nikodym derivative $q^*(\cdot|a)$ is positive; (ii) For all $a \in \mathcal{A}$, $u(a, \cdot) \in L^1(\mathcal{Y}, \mathbb{R}, Q^*(\cdot|a))$.*[2]

The above assumptions are standard in the literature. In the special case where $\mathcal{Y}$ is discrete, $q^*(\cdot|a)$ is simply the probability mass function; when $\mathcal{Y}$ is a continuum, $q^*(\cdot|a)$ is the probability density function. Assumption 1(ii) ensures that the agent's expected period-t payoff, $\overline{u}_t := \int_{\mathcal{Y}} u(a_t, y) q^*(y|a_t) \nu(dy)$, is well-defined.

2.2 Subjective Models

The agent does not know the true DGP; instead, she turns to subjective models to learn about it. A subjective model, indexed by θ, consists of two components: (1) a subjective parameter set Ω^θ and (2) a profile of conditional signal distributions, $Q^\theta : \mathcal{A} \times \Omega^\theta \to \Delta\mathcal{Y}$. One can capture any parameter uncertainty by appropriately specifying a non-singleton Ω^θ. I restrict attention to subjective models that satisfy the following assumption.

Assumption 2. *For all $a \in \mathcal{A}$: (i) Ω^θ is a finite subset of a Euclidean space; (ii) for all $\omega \in \Omega^\theta$, $Q^\theta(\cdot|a, \omega)$ is absolutely continuous w.r.t. the measure ν, and the Radon-Nikodym derivative $q^\theta(\cdot|a, \omega)$ is positive; (iii) for all $\omega \in \Omega^\theta$, $u(a, \cdot) \in L^1\left(\mathcal{Y}, \mathbb{R}, Q^\theta(\cdot|a, \omega)\right)$; (iv) for all $\omega \in \Omega^\theta$, there exists $g_a \in L^2(\mathcal{Y}, \mathbb{R}, \nu)$ such that $\left|\ln \frac{q^*(\cdot|a)}{q^\theta(\cdot|a,\omega)}\right| \leq g_a(\cdot)$ a.s.-$Q^*(\cdot|a)$.*

Assumption 2(i) requires that the parameter space is finite. Assumptions 2(ii) and 2(iii) are analogous to Assumption 1. They ensure that the subjective models do not rule out events that occur with positive probability under the true DGP. Assumption 2(iv) guarantees

[2] $L^p(\mathcal{Y}, \mathbb{R}, \nu)$ denotes the space of all functions $g : \mathcal{Y} \to \mathbb{R}$ s.t. $\int |g(y)|^p \nu(dy) < \infty$.

that the distance between the predictions of the model and the true DGP can be properly quantified so that we can establish a law of large numbers.

Let Θ be the set of all models θ that satisfy Assumption 2. Since each element of Θ is a finite vector of conditional distributions, we have $\Theta \subset \cup_{z=1}^{\infty} (\Delta\mathcal{Y})^{|\mathcal{A}|z}$, where z represents the size of the parameter set. A model θ is said to be *correctly specified* if $q^*(\cdot|a) \equiv q^\theta(\cdot|a,\omega), \forall a \in \mathcal{A}$ for some $\omega \in \Omega^\theta$, i.e. the profile of conditional distributions under θ includes the true DGP, and *misspecified* otherwise.

2.3 The Switcher's Problem

The agent has access to a finite set of subjective models, $\Theta^\dagger \subset \Theta$. It is often assumed in the misspecified learning literature that the decision maker is a *dogmatic modeler* who has a single subjective model, denoted by $\Theta^\dagger = \{\theta\}$. Starting with a full-support prior $\tilde{\pi}_0^\theta \in \Delta\Omega^\theta$, the dogmatic modeler updates her belief based on the history, i.e. $\tilde{\pi}_t^\theta = B^\theta\left(a_{t-1}, y_{t-1}, \tilde{\pi}_{t-1}^\theta\right)$, where $B^\theta : \mathcal{A} \times \mathcal{Y} \times \Delta\Omega^\theta \to \Delta\Omega^\theta$ is the Bayesian operator. The dogmatic modeler then chooses an action to maximize the expected sum of discounted payoffs.

The key departure I take here is to focus on a *switcher* who entertains at least two subjective models. She assigns to each $\theta \in \Theta^\dagger$ a full-support prior $\pi_0^\theta \in \Delta\Omega^\theta$. The agent starts by adopting the *initial model* $\theta^0 \in \Theta^\dagger$, while evaluating a finite set of *competing models* $\Theta^c := \Theta^\dagger \setminus \{\theta^0\}$. Denote as $m_t \in \Theta^\dagger$ the model choice in period t, where $m_0 = \theta^0$. I now describe the events happening in period t in chronological order.

Model switching. The agent employs a *Bayes factor* rule to determine m_t. Fix a constant $\alpha > 1$ that I call the *switching threshold*.[3] At the beginning of each period $t \geq 1$, the agent calculates a vector of Bayes factors $\boldsymbol{\lambda}_t = \left(\lambda_t^\theta\right)_{\theta\in\Theta^\dagger}$, where

$$\lambda_t^\theta = l_t^\theta / l_t^{m_{t-1}}, \tag{1}$$

and

$$l_t^\theta = \sum_{\omega\in\Omega^\theta} \pi_0^\theta(\omega) \prod_{\tau=0}^{t-1} q^\theta(y_\tau|a_\tau,\omega). \tag{2}$$

That is, λ_t^θ is the ratio of the likelihood of model θ to the likelihood of the last period's model choice m_{t-1}. Let $\theta^* := \arg\max_{\theta\in\Theta^\dagger} \lambda_t^\theta$. If $\lambda_t^{\theta^*} > \alpha$, then a switch to θ^* is triggered, i.e. $m_t = \theta^*$; if $\lambda_t^{\theta^*} \leq \alpha$, then the agent does not switch, $m_t = m_{t-1}$. The switching threshold does not change with the direction of switching.[4]

[3] Our analysis in Section 4 goes through if I alternatively assume that the agent repeatedly draws α from a distribution G with $\text{supp}(G) \subset (1,\infty)$.

[4] This symmetry in the switching threshold is made to simplify notation. All results remain unchanged if

Essentially, the agent is conducting a thought experiment: had I adopted an alternative model, would it better explain the observations? As α becomes larger, switching requires stronger evidence. Thus, α can be seen as a measure of the agent's "stubbornness", status quo bias, or a reduced-form indicator of the cost of shifting the paradigm. I discuss the role of a fixed α in Section 2.5.

Learning. After pinning down the model choice, the agent updates her belief under m_t using the full history. For each $\theta \in \Theta^\dagger$, I recursively define a belief process,

$$\pi_t^\theta = B^\theta(a_{t-1}, y_{t-1}, \pi_{t-1}^\theta). \tag{3}$$

However, note that the switcher need not keep track of her posteriors for all models in all periods. Rather, she computes a posterior only for the model that is currently adopted, $\pi_t^{m_t}$, because this is all she needs to make decisions.

Actions. The agent maximizes the sum of discounted expected payoff under her current model m_t. Conditional on adopting θ and holding a belief $\pi^\theta \in \Delta\Omega^\theta$, she solves the following dynamic programming problem,

$$U^\theta\left(\pi^\theta\right) = \max_{a\in\mathcal{A}} \sum_{\omega\in\Omega^\theta} \pi\left(\omega\right) \int_{y\in\mathcal{Y}} \left[u\left(a,y\right) + \delta U^\theta\left(\pi^{\theta\prime}\right)\right] q^\theta\left(y|a,\omega\right) v\left(dy\right), \tag{4}$$

where $\pi^{\theta\prime} = B^\theta\left(a, y, \pi^\theta\right)$. Denote the solution to the above problem as $A^\theta : \Delta\Omega^\theta \rightrightarrows \mathcal{A}$. The agent plays according to a pure optimal policy $a^\theta : \Delta\Omega^\theta \to \mathcal{A}$ such that $a^\theta\left(\pi^\theta\right) \in A^\theta\left(\pi^\theta\right), \forall\pi^\theta \in \Delta\Omega^\theta$. For convenience, let $A_m^\theta : \Delta\Omega^\theta \rightrightarrows \mathcal{A}$ denote the set of myopically optimal actions, which is the solution to (4) when $\delta = 0$. Notice that while experimentation within a model is allowed, the agent does not actively experiment which model is better. I discuss this assumption in Section 2.5.

The underlying probability space $(Y, \mathcal{F}, \mathbb{P})$ is constructed as follows. The sample space is $Y := (\mathcal{Y}^\infty)^{\mathcal{A}}$, each element of which is an infinite sequence of outcome realizations $(y_{a,0}, y_{a,1}, ...)_{a\in\mathcal{A}}$, where $y_{a,\tau}$ denotes the outcome when the agent takes $a \in \mathcal{A}$ in period τ. Denote by $\mathbb{P}$ the probability measure induced by independent draws from q^* and denote by $\mathcal{F}$ the product sigma algebra. Let $h := (a_\tau, y_\tau)_{\tau=0}^\infty$ denote an infinite history and $H := (\mathcal{A} \times \mathcal{Y})^\infty$ be the set of infinite histories. Together with the switching threshold α, the set of models $\Theta^\dagger$, the initial model θ^0, the priors and policies $\left(\pi_0^\theta, a_0^\theta\right)_{\theta\in\Theta^\dagger}$, $\mathbb{P}$ induces a probability measure $\mathbb{P}_S$ on the infinite histories of a switcher. Meanwhile, $\mathbb{P}$ and $\left(\pi_0^\theta, a_0^\theta\right)$ induce a probability measure $\mathbb{P}_B$ on the histories of the dogmatic modeler who believes in θ.

the agent uses different thresholds for different switches.

2.4 Persistence of Models

Fixing an initial model and a set of competing models, the sequence of adopted models (m_t) is a stochastic process that can oscillate forever or converge to one model, and this limit behavior can be path-dependent. We are interested in the situation where the agent eventually settles down with the initial model. For convenience, given any set of models $\Theta' \subset \Theta$, I write the vector of priors $\left(\pi_0^\theta\right)_{\theta\in\Theta'}$ as $\boldsymbol{\pi}_0^{\Theta'}$ and the vector of policies $\left(a^\theta\right)_{\theta\in\Theta'}$ as $\boldsymbol{a}^{\Theta'}$.

Definition 1. Model θ *persists against* $\Theta^c = \Theta^\dagger \backslash \{\theta\}$ (or *persists in* $\Theta^\dagger$) at $\boldsymbol{\pi}_0^{\Theta^\dagger}$ and $\boldsymbol{a}^{\Theta^\dagger}$ if, given a switcher who is endowed with $\Theta^\dagger$, $\boldsymbol{\pi}_0^{\Theta^\dagger}$, $\boldsymbol{a}^{\Theta^\dagger}$, and uses θ as the initial model, the model choice m_t converges to θ with positive probability.

To ease exposition, I will say that θ persists against θ^c if θ persists against $\Theta^c = \{\theta^c\}$.[5] Persistence will play a crucial role in our definition of model robustness. If θ does not persist against Θ^c, at least one different model must be adopted by the switcher infinitely many times. Consequently, the long-term beliefs and behavior of the switcher can be quite different from the predictions of an analyst who only knows the initial model θ and the true DGP. In later sections, I will focus on the following question: what properties should a model possess so that it persists against a feasible set of competing models? The answer to that question depends on not only the set of competing models but also the primitives of the agent, including the prior and policy adopted for each model. Hence, by varying the restrictions imposed on the set of competing models and the primitives, we can obtain different notions of robustness.

2.5 Discussion of Modeling Choices

Before proceeding to the analysis, I briefly comment on several important assumptions of this framework.

Sticky switching. As has been discussed in the introduction, models and parameters are conceptually different, despite their similar roles in determining the outcome distribution. For example, for the theory of classical mechanics, while the gravitational constant is clearly a parameter to be estimated, quantum mechanics is another model that builds on fundamentally different assumptions and has different parameters to estimate; for a deeply overconfident agent, the ability of his coworker is a parameter, while correcting his self-perception is equivalent to shifting to a new model. Model switching features stickiness,

[5]If $\Theta^\dagger$ is not a singleton, then persistence against $\Theta^\dagger$ in general does not imply persistence against each model in $\Theta^\dagger$. See Appendix C for an example.

potentially due to the physical or mental cost of discovering and shifting paradigms. Thus, a switch only occurs when evidence reveals that the alternative is sufficiently better, which is captured by the assumption that $\alpha > 1$.

No experimentation across models. The agent is myopic when it comes to model switching—she chooses her best response presuming that no switch will happen in the future. Thus, she does not actively experiment which model is better, but passively switches according to the Bayes factor rule. Again, this highlights the stickiness of switching models as opposed to the smoothness of Bayesian updating. This assumption is most reasonable when the environment is complex and switching happens rarely (high switching threshold). For instance, consider a scientist who has a model in Newtonian mechanics but is aware of the existence of general relativity. Running experiments to calibrate his model is hard enough, so he may not be spending additional resources to actively experiment and distinguish the two models. However, he will indeed switch to a different model if his experiments turn out to suggest that general relativity is important. Of course, this can be true for ordinary people as well—think of a flat-earth believer who does not actively test if his belief is wrong, but may change his mind after hearing from an old friend.

Comparison with other switching rules. The Bayes factor rule enjoys a few advantages over other common criteria for model selection. First, it has a strong "Bayesian" flavor since the agent does nothing more than keeping track of the relative likelihood ratio of models. Hence, the agent maintains, to some extent, conceptual consistency in belief updating and model switching. Second, Bayes factor is flexible in that it could be easily formulated for any model and any outcome structure without specifying details such as the dimensionality of parameter space, while this information is vital for AIC and BIC. Lastly, Bayes factor automatically includes a penalty for including too much structure into the model and thus helps prevent overfitting. This is manifested in the comparison with the LR test, which evaluates the relative likelihood ratio at the maximum likelihood estimate of the parameter. This gives an advantage to models that fit better in the short term but are more complicated.

3 Berk-Nash Equilibria

Our main question concerns when a model θ persists against a set of competing models Θ^c. By definition, if θ persists against Θ^c, then there exists a first period T such that, with positive probability, the switcher adopts θ and never switches to other models in Θ^c thereafter. From there, her behavior will be identical to a dogmatic modeler who shares the same policy and same posterior for θ at the onset of period T. Hence, whether the switcher can hold on to θ is closely related to how a dogmatic modeler behaves.

Characterizing the asymptotic behavior of a dogmatic modeler is an important question of the misspecified learning literature. A key finding of the literature is that whenever the modeler's behavior stabilizes, the limit behavior must constitute a Berk-Nash equilibrium (Esponda and Pouzo, 2016; Esponda et al., 2019). I now briefly introduce necessary notation to define the Berk-Nash equilibrium and related concepts, including the self-confirming equilibrium (Fudenberg and Levine, 1993) and some refinements. Familiar readers may proceed directly to Section 4.

Denote the Kullback-Leibler divergence (henceforth, KL divergence) of a density q from another density q' as $D_{KL}(q \parallel q')$, where

$$D_{KL}(q \parallel q') := \int_{\mathcal{Y}} q \ln(q/q') \, \nu(dy). \tag{5}$$

The KL divergence of q from q' is an asymmetric non-negative distance measure between q and q', which is minimized to zero if and only if q and q' coincide almost everywhere. With a slight abuse of notation, given any strategy σ, let

$$\Omega^\theta(\sigma) := \arg\min_{\omega' \in \Omega^\theta} \sum_{\mathcal{A}} \sigma(a) \, D_{KL}\left(q^*(\cdot|a) \parallel q^\theta(\cdot|a, \omega')\right). \tag{6}$$

That is, $\Omega^\theta(\sigma) \subseteq \Omega^\theta$ identifies the *KL minimizers at σ under θ*, i.e. the parameters in Ω^θ which yield the closest match to the true DGP when the agent plays σ.

Definition 2. Strategy $\sigma \in \Delta\mathcal{A}$ is a *Berk-Nash equilibrium* (BN-E) under θ if there exists a belief $\pi \in \Delta\Omega^\theta(\sigma)$ with $\sigma \in \Delta A_m^\theta(\pi)$. A BN-E σ is

(i) *quasi-strict* if there exists a belief $\pi \in \Delta\Omega^\theta(\sigma)$ with $\text{supp}(\sigma) = A_m^\theta(\pi)$.

(ii) *uniformly quasi-strict* if $\text{supp}(\sigma) = A_m^\theta(\pi)$ for every belief $\pi \in \Delta\Omega^\theta(\sigma)$.

(iii) a *self-confirming equilibrium* (SCE) if there exists a belief $\pi \in \Delta\Omega^\theta(\sigma)$ with $\sigma \in \Delta A_m^\theta(\pi)$ and $q^*(\cdot|a) \equiv q^\theta(\cdot|a, \omega), \forall \omega \in \text{supp}(\pi), \forall a \in \text{supp}(\sigma)$.

A Berk-Nash equilibrium requires myopic optimality against a belief π that takes support on the KL minimizers; both σ and π could be non-degenerate. Quasi-strictness further requires that σ takes support on all myopically optimal actions at π. A *strict* BN-E is a pure quasi-strict BN-E. Uniformly quasi-strictness is stronger than quasi-strictness for it additionally requires that the set of myopically optimal actions remains unchanged as long as the belief takes support on the KL minimizers at σ. Further, a BN-E is *uniformly strict* if it is pure and uniformly quasi-strict.[6] Finally, a self-confirming equilibrium requires that each parameter in the support of π induces the same outcome distribution as the true DGP does at each action a that is played with positive probability. A *quasi-strict SCE* and a *uniformly quasi-strict SCE* can be defined analogously. While every subjective model admits at least one Berk-Nash equilibria (Esponda and Pouzo, 2016), the existence of a self-confirming equilibrium is not guaranteed.

4 Robustness

4.1 Global Robustness

Our notions of robustness are defined based on the long-term persistence of a model. I first define *global robustness*, which requires a model to persist against all possible sets of competing models.

Definition 3 (Global robustness)**.** A model θ is *globally robust* if for every finite $\Theta^c \subset \Theta$, there exists a full-support π_0^θ and an optimal a^θ, under which θ persists against Θ^c at all full-support $\boldsymbol{\pi}_0^{\Theta^c}$ and optimal $\boldsymbol{a}^{\Theta^c}$.

The interpretation of global robustness is as follows. Global robustness guarantees that θ can persist no matter what alternative models it is compared against, as long as the agent starts with a proper choice of prior and policy for θ. Note that this choice can vary with the set of competing models. Conversely, if θ is not robust, then one can find a set of competing models associated with some prior and some policy so that at least one competing model is almost surely (a.s.) adopted infinitely often, regardless of the prior and policy assigned to θ.

Given a particular Θ^c, figuring out which model persists against Θ^c can be quite challenging. However, globally robust models are not as hard to characterize, thanks to the requirement that they can persist against *every* possible $\Theta^c \subset \Theta$. In particular, let $\Theta^c = \{\theta^c\}$ and assume for now that θ^c is correctly specified. Since θ^c is correctly specified, the likelihood ratio of θ to θ^c, or $l_t^\theta / l_t^{\theta^c}$, is a martingale that a.s. converges.[7] Hence, on paths where θ is eventually forever adopted, the inverse likelihood ratio $l_t^{\theta^c}/l_t^\theta$—which eventually equals $\lambda_t^{\theta^c}$—a.s. converges to some value below α. Furthermore, decomposing the posteriors π_t^θ and $\pi_t^{\theta^c}$ by a few different likelihood ratios, I show that π_t^θ and $\pi_t^{\theta^c}$ also a.s. converge on those paths. We thus obtain the following lemma.

[6]Fudenberg et al. (2021) first define a uniformly strict BN-E.
[7]See Lemma 3 in Appendix A.

Lemma 1. *Suppose a model $\theta \in \Theta$ persists against a correctly specified model $\theta^c \in \Theta$ at some full-support $\pi_0^\theta, \pi_0^{\theta^c}$ and optimal a^θ, a^{θ^c}. Then on paths where m_t converges to θ, almost surely, $l_t^{\theta^c}/l_t^\theta$ converges to a random variable $\iota \leq \alpha$, π_t^θ converges to a random variable $\pi_\infty^\theta \in \Delta\Omega^\theta$, and $\pi_t^{\theta^c}$ converges to a random variable $\pi_\infty^{\theta^c} \in \Delta\Omega^{\theta^c}$.*

All proofs of the main results are in Appendix B. An implication of Lemma 1 is that θ must be able to perfectly predict the distribution of outcomes in the long term. This observation follows from the fact that with a correctly specified model, a learner a.s. assigns probability close to 1 to the true outcome distribution in the limit (Easley and Kiefer, 1988). Suppose θ^c persistently outperforms θ in explaining the observed outcomes, then by the Law of Large Numbers, the likelihood ratio $l_t^{\theta^c}/l_t^\theta$ a.s. grows to infinity. Hence, π_t^θ must also assign probability close to 1 to the true outcome distribution in the limit on a positive measure of paths.

Since data is endogenously generated, this further implies that the agent ends up playing a self-confirming equilibrium with positive probability. More precisely, since $l_t^{\theta^c}/l_t^\theta$ would perpetually fluctuate if the agent plays non-equilibrium actions infinitely often, the agent should end up playing only the equilibrium actions with positive probability. Our next definition formalizes this property. On paths where θ is adopted forever, a switcher eventually behaves no differently than a dogmatic modeler. This condition can therefore be cast as a property of the dogmatic modeler.[8]

Definition 4. A BN-E $\sigma \in \Delta\mathcal{A}$ under θ is said to be *absorbing with positive probability*, or *p-absorbing* if under some full-support π_0^θ and optimal a^θ, there exists $T \geq 0$ such that, with positive probability, a dogmatic modeler with $\Theta^\dagger = \{\theta\}$ only plays actions in $\text{supp}\,(\sigma)$ after period T.

That a BN-E with support $A \subset \mathcal{A}$ is p-absorbing does not imply that the dogmatic modeler's action process converges to a single action in A or her action frequency converges to a certain mixed strategy with positive probability.[9] Rather, it allows for non-convergent behavior within A but rules out the scenario where the modeler a.s. plays actions outside A infinitely often. I conclude our analysis for the case of a correctly specified competing model with the following lemma.

Lemma 2. *Suppose a model $\theta \in \Theta$ persists against a correctly specified model $\theta^c \in \Theta$ at some full-support $\pi_0^\theta, \pi_0^{\theta^c}$ and optimal a^θ, a^{θ^c}. Then there exists a p-absorbing SCE under θ.*

[8] I assume the dogmatic modeler and the switcher have the same discount factor throughout.

[9] For example, this is weaker than the stability notion proposed by Fudenberg et al. (2021). They define that a pure BN-E a^* under θ is stable if for every $\kappa \in (0,1)$, there exists a belief $\pi \in \Delta\Omega^\theta$ and some $\epsilon > 0$ such that for any prior $\pi_0^\theta \in B_\epsilon(\pi)$, the dogmatic modeler's action sequence a_t converges to a^* with probability larger than κ.

It now becomes clear that persisting against a correctly specified model conveys abundant information about θ. Perhaps surprisingly, this alone is powerful enough to guarantee that the model also persists against every other finite set of competing models. Theorem 1 shows that the existence of a p-absorbing SCE is not only necessary but also sufficient for global robustness.

Theorem 1. *A model $\theta \in \Theta$ is globally robust if and only if there exists a p-absorbing SCE under θ.*

The seemingly demanding notion of global robustness amounts to the requirement that θ persists against one arbitrary correctly specified model. For instance, provided that θ can beat a competing model that assigns a tiny probability to the true DGP, it also has the potential to beat one that assigns probability 1 to the true DGP. Conversely, models that fail to be globally robust will not persist in the long term as long as the agent evaluates some correctly specified model. More importantly, Theorem 1 reveals the equivalence between global robustness and the existence of a p-absorbing self-confirming equilibrium under θ, a property that can be further characterized using tools from the existing literature since it only concerns the problem of a dogmatic modeler. It thus provides a foundation for the persistence of certain types of misspecification.

I now briefly outline why the existence of a p-absorbing SCE implies global robustness. Suppose σ is a p-absorbing SCE under θ, then under some prior π_0^θ and policy a^θ, with positive probability, a dogmatic modeler's action converges to the support of σ and, without loss of generality, each action in the support of σ is played infinitely often. When such convergence occurs, the dogmatic modeler's belief a.s. converges and the limit belief assigns probability 1 to the KL minimizers $\Omega^\theta(\sigma)$. Since σ is self-confirming, all KL minimizers predict an outcome distribution that is identical to the true DGP at all actions in the support of σ. Since no models can consistently outperform the true DGP in matching the data, if we evaluate the fitness of competing models on these histories, the likelihood ratio of θ to each competing model θ^c, $l_t^\theta / l_t^{\theta^c}$, is asymptotically bounded below by the probability that π_0^θ assigns to the KL minimizers $\Omega^\theta(\sigma)$. Taken together, these observations imply that we can find a new prior $\hat{\pi}_0^\theta$ that is sufficiently concentrated on $\Omega^\theta(\sigma)$, such that on a positive measure of histories $\hat{H}$, this new prior induces the dogmatic modeler to only play actions in $\operatorname{supp}(\sigma)$ and every Bayes factor computed from $\hat{H}$ never exceeds α. It follows that a switcher and a dogmatic modeler must behave identically on $\hat{H}$, since no switch would ever happen.[10] Therefore, if the switcher starts with the same prior $\hat{\pi}_0^\theta$ and the same policy a^θ, it is a positive-probability event that she will only play actions from $\operatorname{supp}(\sigma)$ and adopt θ

[10]This is akin to the coupling argument commonly used in probability theory.

forever.

An important feature of Theorem 1 is that the characterization of global robustness does not depend on the switching threshold α. However, the proof of Theorem 1 shows that α is indeed relevant to the choice of the prior. As α decreases, model switching becomes less sticky. Thus, the agent's prior needs to be more entrenched and concentrated around KL minimizers such that she finds her initial model good enough compared to any competing model.

An immediate corollary of Lemma 2 and Theorem 1 is that any correctly specified model is globally robust since a model must persist against itself.[11]

Corollary 1. *Every correctly specified model is globally robust.*

Corollary 2 takes a different route and provides a sufficient condition for an SCE to be p-absorbing, which can be easily verified from the primitives. In contrast to Corollary 1, this corollary shows that misspecified models can be globally robust.

Corollary 2. *A model $\theta \in \Theta$ is globally robust if there exists a uniformly quasi-strict SCE under θ.*

The proof of Corollary 2 is similar to Theorem 2 in Fudenberg et al. (2021) and Theorem 1 in Frick et al. (2020). Suppose we have a uniformly quasi-strict SCE σ with a supporting belief $\pi \in \Delta\Omega^{\theta}$. First, since σ maximizes the flow payoff against every $\omega \in \text{supp}\,(\pi)$, there is no experimentation incentive to distinguish parameters in $\text{supp}\,(\pi)$; thus, it must be that the set of dynamically optimal actions at π coincides with the myopically optimal actions. Then by the upper-hemicontinuity of A^{θ}, we can find a small open ball around π, denoted by $B_{\epsilon}(\pi) \subset \Delta\Omega^{\theta}$, such that at any belief $\pi' \in B_{\epsilon}(\pi)$, the optimal action(s) $A^{\theta}(\pi')$ are contained in the support of σ. Finally, since the equilibrium is self-confirming, we invoke Ville's maximal inequality for supermartingales (Ville, 1939) to show that as long as a dogmatic modeler's prior is close enough to π, with positive probability, her posterior never leaves the neighborhood $B_{\epsilon}(\pi)$ and thus she ends up always playing actions in $\text{supp}\,(\sigma)$.

Example 1 demonstrates how overconfidence in one's ability (Heidhues et al., 2019) can give rise to a globally robust misspecified model.

Example 1 (Overconfidence)**.** Consider a discrete version of Example 2 in Heidhues et al. (2019). A worker chooses a level of costly effort each period from $\mathcal{A} = \{0, 1, 2\}$ and observes a payoff of $u(a_t, y_t) = y_t$. The true DGP determines the output, $y_t = (a_t + b^*)\,\omega^* - .5a_t^2 + \eta_t$, where $b^* = 1$ is his true ability level, $\omega^* = 2$ is an environment fundamental that determines

[11] Notice that the likelihood ratio between one model and itself is always 1.

the return to effort, and η_t is a zero-mean random noise distributed according to f. The efficient effort level is $a^* = 2$. The worker is initially overconfident in that he believes his ability is given by $\tilde{b} = 3 > b^*$. He is unsure of the return to effort and needs to learn about it over time. This is reflected from his initial subjective model θ, in which he treats the return to effort as a parameter to be estimated: according to θ, the output is given by $y_t = \left(a_t + \tilde{b}\right)\omega - .5a_t^2 + \eta_t$, where ω is an element of a finite parameter set $\Omega^\theta = \{1, 2\}$ and η_t is distributed according to f.

It can be readily verified that there exists a unique Berk-Nash equilibrium under θ, in which the worker chooses $\tilde{a} = 1$ and believes in $\tilde{\omega} = 1$; due to overconfidence, he attributes his underperformance to a bad environment and exerts lower effort in response. This BN-E is also a uniformly quasi-strict SCE. Thus, θ must be globally robust. The interpretation of this result is as follows. Suppose the worker has had a performance review and starts to evaluate a competing model θ^c that is correctly specified about the ability. That is, θ^c predicts that $y_t = (a_t + b^*)\,\omega - .5a_t^2 + \eta_t$, where $\Omega^{\theta^c} = \Omega^\theta = \{1, 2\}$. Nevertheless, it turns out that this competing model does not appear a lot more compelling when the worker keeps exerting low effort and has a prior π_0^θ that is sufficiently concentrated on $\tilde{\omega} = 1$. As a consequence, the worker, who prefers the status quo, can be perpetually trapped in the inefficient state of being overconfident and choosing low effort.

4.2 Unconstrained Local Robustness

It may be implausible for the agent to evaluate a correctly specified model or a competing model that considerably differs from the initial model, especially when the environment is complex or when the agent is only willing to consider small changes. Although global robustness has a clean characterization, global robustness may not be the best criterion under those circumstances to evaluate if some form of misspecification is likely to persist. A natural question is then whether a model persists against *neighbor* models. When the answer is no, then such a model cannot be adopted forever even when the agent evaluates slightly different models. By contrast, if the answer is yes, then this model can persist whenever the switcher only takes small steps. Following this line of thought, I develop a weaker robustness property, *local robustness.*

First, I formalize what qualifies as a neighbor model. In Sections 4.2 and 4.3, I probe into two different approaches to defining neighbor models, non-parametric and parametric, and show that they give rise to two distinct notions of local robustness. Since every model consists of nothing more than a profile of action-contingent outcome distributions, a direct measure of proximity of models is the distance between the sets of distributions. In this subsection, I develop a notion of local robustness based on this measure, which is nonparametric, and therefore, *unconstrained*. In contrast, in the next subsection, I introduce a parametric approach that places constraints on the structure of neighbor models.

Consider two action-contingent outcome distributions Q and Q', both of which are elements of $(\Delta\mathcal{Y})^{\mathcal{A}}$. I define their distance as the maximum Prokhorov distance across actions,

$$d\left(Q,Q'\right) \coloneqq \max_{a\in\mathcal{A}} d_P\left(Q,Q'\right), \tag{7}$$

where

$$d_P\left(Q,Q'\right) = \inf\left\{\epsilon > 0 | Q_a\left(Y\right) \leq Q_a\left(B_\epsilon\left(Y\right)\right) + \epsilon \text{ for all } Y \subset \mathcal{Y}\right\} \tag{8}$$

denotes the Prokhorov metric. For any set of action-contingent outcome distributions, $\mathcal{Q} \subseteq (\Delta\mathcal{Y})^{\mathcal{A}}$, define $\mathcal{Q}_\epsilon \coloneqq \left\{Q' \in (\Delta\mathcal{Y})^{\mathcal{A}} : \exists Q \in \mathcal{Q}, \text{s.t. } d\left(Q,Q'\right) \leq \epsilon\right\}$, which is the ϵ-ball around $\mathcal{Q}$.

Now we are ready to define neighbor models. Let $Q^{\theta,\omega} \coloneqq \left\{Q^\theta\left(\cdot|a,\omega\right)\right\}_{a\in\mathcal{A}}$ denote the action-contingent outcome distribution induced by $\omega \in \Omega^\theta$ under model θ, and let $\mathcal{Q}^\theta \coloneqq \left\{Q^{\theta,\omega}\right\}_{\omega\in\Omega^\theta}$ denote all distributions in the support of θ. Define the ϵ-neighborhood of θ as

$$N_\epsilon\left(\theta\right) \coloneqq \left\{\theta' \in \Theta : d_H\left(\mathcal{Q}^\theta, \mathcal{Q}^{\theta'}\right) \leq \epsilon\right\}, \tag{9}$$

where

$$d_H\left(\mathcal{Q}^\theta, \mathcal{Q}^{\theta'}\right) \coloneqq \max\left\{\max_{\omega\in\Omega^\theta}\min_{\omega'\in\Omega^{\theta'}} d\left(Q^{\theta,\omega}, Q^{\theta',\omega'}\right), \max_{\omega'\in\Omega^{\theta'}}\min_{\omega\in\Omega^\theta} d\left(Q^{\theta,\omega}, Q^{\theta',\omega'}\right)\right\} \tag{10}$$

denotes the Hausdorff metric. Notice in particular there is no restriction relating the parametric family of θ with θ'. I now define unconstrained local robustness.

Definition 5. A model $\theta \in \Theta$ is *unconstrained locally robust* if there exists some $\epsilon > 0$, such that for every finite $\Theta^c \subset N_\epsilon\left(\theta\right)$, there exists a full-support π_0^θ and an optimal a^θ, under which θ persists against Θ^c at all full-support $\boldsymbol{\pi}_0^{\Theta^c}$ and optimal $\boldsymbol{a}^{\Theta^c}$.

The definition of local robustness seems much weaker than global robustness. It only requires the model to be able to persist against neighbor models. When θ is misspecified, a sufficiently nearby model is necessarily also misspecified. This prevents us from invoking Lemmas 1 and 2 and inferring that a p-absorbing SCE must exist. However, Theorem 2 shows that these two notions of robustness point to the same set of subjective models.

Theorem 2. *Unconstrained local robustness is equivalent to global robustness.*

The idea of Theorem 2 is quite simple. Given any unconstrained locally robust model θ, we could construct a neighbor competing model θ^c with the same parameter set and the same prior but potentially different outcome distributions, such that each parameter corresponds to a distribution that is a convex combination of the true DGP and the corresponding distribution under θ. Then the likelihood ratio of θ^c to θ would also be a linear combination of 1 and the likelihood ratio of the true DGP to θ. Hence, θ persists against θ^c only if θ also persists against the true DGP, which makes it globally robust.

Theorem 2 also provides a new perspective for understanding global robustness. If model θ fails to be globally robust, the switcher need not go far to find an attractive alternative—taking small but undirected steps is as powerful as taking big steps.

4.3 Constrained Local Robustness

The notion of local robustness introduced in the previous section has not restricted the nature of models that are being compared by the switcher. But in many cases when individuals are evaluating competing models, they are perturbing their initial model while maintaining the same parametric structure. I now take a parametric approach by restricting attention to a profile of conditional outcome densities $\{p(\cdot|a,\omega)\}_{a\in\mathcal{A},\omega\in\Omega^p}$ that are uniformly continuous in $\omega \in \Omega^p \subseteq \mathbb{R}^d$ for all $a \in \mathcal{A}$, where Ω^p is the parameter set associated with p.

Definition 6. The *p-family* of models is the set $\Theta^p \subset \Theta$ such that

$$\Theta^p \coloneqq \{\theta \in \Theta : q^\theta(\cdot|a,\omega) \equiv p(\cdot|a,\omega) \text{ for all } \omega \in \Omega^\theta \subseteq \Omega^p \text{ and all } a \in \mathcal{A}\}. \tag{11}$$

Two models θ and θ' that belong to the same family Θ^p share the same mapping from parameters to outcome distributions but differ in their parameter sets. We can then conveniently measure their distance by the Hausdorff distance between Ω^θ and $\Omega^{\theta'}$. Formally, define an ϵ-neighborhood of θ in Θ^p as

$$N_\epsilon^p(\theta) \coloneqq \left\{\theta' \in \Theta^p : d_H\left(\Omega^\theta, \Omega^{\theta'}\right) \leq \epsilon\right\}, \tag{12}$$

where $d_H\left(\Omega^\theta, \Omega^{\theta'}\right) = \max\left\{\max_{\omega\in\Omega^\theta}\min_{\omega'\in\Omega^{\theta'}}\|\omega-\omega'\|, \max_{\omega'\in\Omega^{\theta'}}\min_{\omega\in\Omega^\theta}\|\omega-\omega'\|\right\}$. Now we can define constrained local robustness.

Definition 7. A model $\theta \in \Theta^p$ is *p-constrained locally robust* if there exists some $\epsilon > 0$, such that for every finite $\Theta^c \subset N_\epsilon^p(\theta)$, there exists a full-support π_0^θ and an optimal a^θ, under which θ persists against Θ^c at all full-support $\boldsymbol{\pi}_0^{\Theta^c}$ and optimal $\boldsymbol{a}^{\Theta^c}$.

The notion of p-constrained local robustness requires a model to be able to persist against all nearby models within the p-family. Below I present two examples that fit into the parametric setting.

Example 2 (Underestimation of financial risks)**.** The agent is an investor who chooses her investment level $a_t \in \mathcal{A} = \{1, 2\}$. In each period, she obtains a flow utility of $u(a_t, y_t) = 1 - e^{-y_t}$, where $y_t \in \mathcal{Y} = \mathbb{R}$ is the investment return. Notice that the agent is risk-averse as her utility is concave in y_t. When the investor chooses a_t, the true return is given by $y_t = 2a_t - \frac{1}{2}a_t^2 + \xi_t$, where ξ_t is a zero-mean normally distributed variable with variance a_t, i.e. $\xi_t \sim N(0, a_t)$. Playing a higher action not only changes the mean of the return, but also increases its variance. The agent initially posits that $y = \omega_1 a_t - \frac{1}{2}a_t^2 + \xi_t$, where ω_1 is an element of $\Omega^\theta = \{1, 2\}$ and $\xi_t \sim N(0, v)$ for some constant $v > 0$. That is, the investor neglects how the investment level affects her risk exposure.

There are different ways to embed θ into a parametric family. First, let $p(\cdot|a, \omega_1)$ be the same normal density as the one predicted by ω_1 under θ, and let $\Omega^p = \mathbb{R}_+$. Then by evaluating a neighbor model in Θ^p, the investor effectively considers a different expected return function. Second, whereas it seems that the parameter space is one-dimensional, we can augment it by adding a second dimension that parameterizes the variance of ξ. This augmentation leads to a two-dimensional parameter space, $\Omega^{\theta'} = \{1, 2\} \times \{v\}$. Let $p'(\cdot|a, \omega)$ be the normal density function associated with mean $\omega_1 a - \frac{1}{2}a^2$ and variance ω_2, and let the corresponding parameter space be $\Omega^{p'} = \mathbb{R}_+^2$. The augmented model θ' belongs to the p'-family. Consider a competing model $\theta^c \in \Theta^{p'}$ with $\Omega^{\theta^c} = \{1, 2\} \times \{v + \epsilon\}$, whose distance from θ' is exactly ϵ. Let $\Theta^c = \{\theta^c\}$. By evaluating Θ^c, the investor is assessing whether the variance takes a different value; however, she still fails to realize its dependence on her investment level.

Example 3 (Monopolist with a misspecified prior (Nyarko, 1991))**.** The agent is a monopolist who chooses a price $a_t \in \{2, 10\}$ each period and observes the market demand $y_t \in \mathbb{R}$ each period. The true DGP is described by $y_t = 40 - 5a_t + \eta_t$, where η_t is a zero-mean random noise i.i.d. distributed according to density f. The monopolist obtains a flow payoff of $u(a_t, y_t) = a_t y_t$. His initial subjective model θ belongs to the p-family, where p predicts that $y_t = \omega_1 - \omega_2 a_t + \eta_t$, where $(\omega_1, \omega_2) \in \Omega^p = \mathbb{R}^2$. The parameter space Ω^θ is a finite subset of $[12, 32] \times [1, 3]$. Notice that θ is misspecified as the monopolist systematically underestimates both the intercept and the slope of the demand function. Let $\Theta^c = \{\theta^c\}$, where the competing model θ^c belongs to the same family Θ^p but has an expanded parameter space. In particular, let $\Omega^{\theta^c} = \Omega^\theta \cup \left\{\omega' \in \Omega^p : \omega = \omega + (\epsilon, \epsilon) \text{ for some } \omega \in \Omega^\theta\right\}$ for some small $\epsilon > 0$. That is, the monopolist now evaluates a possibility that the parameter may take slightly higher values. It is straightforward to verify that $\Theta^c \subset N_\epsilon^p(\theta)$.

4.3.1 Sufficient Conditions

In this section, I provide sufficient conditions for p-constrained local robustness. Several complicating issues arise as a result of the constraints. First of all, similar to the unconstrained local robustness case, models neighboring a misspecified model must also be misspecified, preventing us from using Lemma 2. Moreover, because of the additional constraints over the parametric structure, it is infeasible to perturb the predictions of model θ unanimously towards the direction of the true DGP. Hence, the set of local robust models can be much larger than the set of globally robust models.

There seems to be an easy fix—similar to Theorem 1, we only need to verify if there exists a p-absorbing Berk-Nash equilibrium σ under θ such that when the agent plays according to σ, no neighbor model is expected to fit the data strictly better than θ. This would be satisfied if all KL minimizers $\Omega^{\theta}(\sigma)$ also locally minimize the KL divergence in the expanded domain Ω^p. Given a family of densities p, define a function $K^p : \Delta\mathcal{A} \times \Omega^p \to \mathbb{R}$, where

$$K^p(\sigma,\omega) := \sum_{\mathcal{A}} \sigma(a)\, D_{KL}\left(q^*(\cdot|a) \parallel p(\cdot|a,\omega)\right). \tag{13}$$

That is, $K^p(\sigma,\omega)$ represents the σ-weighted KL divergence between the outcome distribution predicted by ω and the true DGP. I now define the desired property.

Definition 8. Model θ is *locally KL-minimizing* at σ w.r.t. Ω^p if there exists some $\eta > 0$ such that for all $\omega \in \Omega^{\theta}(\sigma)$ and $\omega' \in B_{\eta}\left(\Omega^{\theta}(\sigma)\right) \cap \Omega^p$, we have $K^p(\sigma,\omega) \le K^p(\sigma,\omega')$.

However, it turns out that we also need a local identification property, as defined below.

Definition 9. Model θ is *locally identified* at σ w.r.t. Ω^p if there exists some $\eta > 0$ such that for all $\omega \in \Omega^{\theta}(\sigma)$ and $\omega' \in B_{\eta}\left(\Omega^{\theta}(\sigma)\right) \cap \Omega^p$, either $K^p(\sigma,\omega) \neq K^p(\sigma,\omega')$ or $p(\cdot|a,\omega') \equiv p(\cdot|a,\omega)$ for all $a \in \mathrm{supp}(\sigma)$.

Suppose θ is not locally identified at σ w.r.t. Ω^p, then one can find a parameter in an arbitrarily small neighborhood of Ω^{θ}, such that it yields the same KL divergence as $\Omega^{\theta}(\sigma)$ yet predicts a different distribution at σ. If so, the likelihood ratio of a nearby model that contains such a parameter resembles a random walk and exceeds the switching threshold infinitely often.[12] Local identification also implies that each parameter in $\Omega^{\theta}(\sigma)$ corresponds to the same outcome distribution at σ.[13] Notice that local identification is automatically satisfied when σ is a self-confirming equilibrium.

[12]Example 7 in Appendix C illustrates why local identification is important.

[13]I take the same view as Esponda and Pouzo (2016) that what really matters for the agent is the distribution of distributions, not individual parameters. Esponda and Pouzo (2016, p. 1108) define a similar concept called weak identification at σ, which requires that two distributions that are judged to be equally a best fit at σ are identical.

I now state a sufficient condition for local robustness. It requires the existence of a p-absorbing Berk-Nash equilibrium, similar to Theorem 1, but also requires local KL-minimization and local identification at the equilibrium. The assumption of the equilibrium being pure is not essential but to ease exposition.[14]

Theorem 3. *A model $\theta \in \Theta^p$ is p-constrained locally robust if there exists a pure p-absorbing BN-E σ under θ such that θ is locally KL-minimizing and locally identified at σ w.r.t. Ω^p.*

In the proof of Theorem 3, I construct a neighborhood of θ in which any competing model θ^c yields either the same distribution as $\Omega^\theta(\sigma)$ does or strictly higher KL divergence at σ. The rest of the proof is similar to the proof of Theorem 1. Analogous to Corollary 2, the following corollary provides a sufficient condition for Theorem 3. As an implication of Theorem 2 in Fudenberg et al. (2021), given any uniformly strict BN-E, there exists a prior belief such that, with positive probability, a dogmatic modeler's action process converges to this equilibrium. This establishes the p-absorbing condition.

Corollary 3. *A model $\theta \in \Theta^p$ is p-constrained locally robust if there exists a uniformly strict BN-E σ under θ such that θ is locally KL-minimizing and locally identified at σ w.r.t. Ω^p.*

Let's revisit Example 2 to illustrate how to apply the results.

Example 2, cont. Had the investor known the true DGP, she would optimally choose the low investment level. However, as the investor neglects the role of more investments in increasing the volatility of her payoff, she always plays the high action since it maximizes the expected return. The only BN-E under both θ and the augmented θ' is the pure action $a = 2$, supported by the belief that assigns probability 1 to $\omega_1 = 2$. By Corollary 3, θ is p-constrained locally robust, as the equilibrium is uniformly strict and the parameter value $\omega_1 = 2$ perfectly matches the mean of the return, which is the only thing the agent is seeking to match.

4.3.2 Necessary Conditions

Restricting the set of competing models and priors also poses a challenge to finding necessary conditions for local robustness. When the competing model is correctly specified and the initial model is eventually adopted, the agent's action process must converge to the support of a Berk-Nash equilibrium. This greatly simplifies the characterization of persistence because we could simply compare θ with the competing model at the BN-E. This convenient

[14] A more general version of Theorem 3 would require the existence of a p-absorbing mixed BN-E σ such that at every $\tilde{\sigma} \in \Delta \operatorname{supp}(\sigma)$, θ is locally KL-minimizing and locally identified w.r.t. Ω^p at $\tilde{\sigma}$. This condition automatically holds when σ is pure or self-confirming.

convergence property is lost when we shift our focus to a competing model that is potentially misspecified. However, as I show in Theorem 4, suppose the action frequency of a dogmatic-modeler with $\Theta^{\dagger} = \{\theta\}$ indeed a.s. converges to some BN-E under every full-support prior and policy, then it is necessary for a p-constrained locally robust model to admit at least one BN-E σ at which θ is locally KL-minimizing w.r.t. Ω^{p}.[15]

Theorem 4. *Suppose that the action frequency of a dogmatic modeler a.s. converges to a BN-E under all full-support priors and policies. Then a model $\theta \in \Theta^{p}$ is p-constrained locally robust only if there exists a BN-E σ under θ such that θ is locally KL-minimizing at σ w.r.t. Ω^{p}.*

Esponda et al. (2019) establish global almost-sure convergence of a dogmatic modeler's action frequency to a "globally attracting" BN-E if it exists and the dogmatic modeler is myopic ($\delta = 0$). The global attractiveness is defined based on a differential equation that describes the evolution of the action frequency. Such convergence can also be observed in a few examples in the literature, all of which impose specific assumptions over the types of misspecification and the outcome distributions (Nyarko (1991); Heidhues et al. (2018); He (2020); Ba and Gindin (2021)). In those environments, Theorem 4 provides a simple criterion to determine if some subjective model is constrained locally robust.

The proof idea of Theorem 4 is as follows. Suppose none of the Berk-Nash equilibria under θ satisfy the local KL-minimization property w.r.t. Ω^{p}. Then for each equilibrium σ, there exists a parameter $\omega' \in \Omega^{p}\backslash\Omega^{\theta}$ that is close to some $\omega \in \Omega^{\theta}$ and an open neighborhood of σ, $B_{\epsilon}(\sigma) \subset \Delta\mathcal{A}$, such that at any strategy $\sigma' \in B_{\epsilon}(\sigma)$, the parameter ω' yields a strictly lower KL divergence than the lowest possible KL divergence under θ. Since the set of Berk-Nash equilibria is compact (Lemma 7), the Heine-Borel theorem implies that there exists a finite set of parameters $\Omega' \subset \Omega^{p}\backslash\Omega^{\theta}$ such that for each equilibrium σ, there exists an $\omega' \in \Omega'$ that satisfies this property. Consider the competing model $\theta^{c} \in \Theta^{p}$ with a larger parameter space $\Omega^{\theta^{c}} = \Omega^{\theta} \cup \Omega'$. Suppose the model choice of the agent converges to θ with positive probability. On the set of paths where she eventually adopts θ forever, she behaves identically to a dogmatic modeler after the final switch to θ, and thus, by assumption, her action frequency converges to a Berk-Nash equilibrium. Under such convergence, θ^{c} strictly outperforms the fit of θ in the long term. Therefore, by the Law of Large Numbers, $\lambda_{t}^{\theta^{c}}$ eventually exceeds α and the switcher adopts θ^{c}. This is a contradiction. Theorem 4 follows.

It is challenging to identify necessary conditions without global convergence of a dogmatic modeler's behavior. Nevertheless, I show in Theorem 5 that the local KL-minimization condition described in Theorem 4 is still necessary when $\mathcal{A}$ is binary and the agent is myopic.

[15]Formally, given a finite action space $\mathcal{A}$ and an action sequence $(a_1, a_2, ...)$, we can construct the action frequency sequence $(\sigma_t)_t$ where $\sigma_t \in \Delta\mathcal{A}$ and $\sigma_t(a) = \frac{1}{t}\sum_{\tau=1}^{t} 1_{\{a_t = a\}}$.

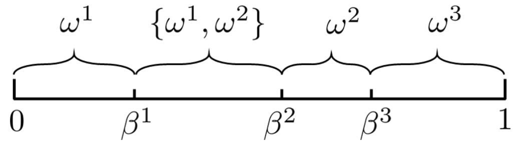

Figure 1: Example of a binary-action setting: Each point in the interval represents a mixed action's weight assigned to a^2; the parameter(s) placed above a segment of the interval are the minimizer(s) of $K^p(\sigma, \omega)$ in Ω^θ for all σ in this segment.

Theorem 5. *Suppose $|\mathcal{A}| = 2$ and $\delta = 0$. Then a model $\theta \in \Theta^p$ is p-constrained locally robust only if there is a BN-E σ under θ such that θ is locally KL-minimizing at σ w.r.t. Ω^p.*

The critical step in proving Theorem 5 is to show that a dogmatic modeler's action frequency almost surely enters an arbitrarily small neighborhood of the set of Berk-Nash equilibria infinitely often. From here, we can use an analogous argument to the proof of Theorem 4. I use Figure 1 to illustrate this first step.

When the action space is binary, we can write any mixed action as $\beta \cdot a^1 + (1-\beta) \cdot a^2$, where $\beta \in [0, 1]$. Therefore, the strategy space can be represented as the unit interval denoting the set of possible weights on a^2. To add more structure, suppose that the parameter space Ω^θ has four elements, each of which is a KL minimizer in Ω^θ at some mixed strategies. Since the KL divergence is continuous in the probability of each action, it is straightforward to show that the set of mixed strategies at which a parameter is a KL minimizer is compact and connected. For example, in Figure 1, ω^1 uniquely minimizes the KL divergence when evaluated at a mixed action when $\beta \in [0, \beta^1]$, while both ω^1 and ω^2 are minimizers when $\beta \in [\beta^1, \beta^2]$. Restrict attention to the set of paths where the sequence of the action frequency $\{\sigma_t\}_t$ is such that both ω^1 and ω^2 are KL minimizers infinitely often but not ω^3. Since the action space is binary, if σ_t enters two non-connected regions on the unit interval infinitely often, it must also cross the region in between infinitely often.[16] This implies that σ_t must enter $[\beta^1, \beta^2]$ infinitely often. To generate this pattern, it must be that $a^2 \in A^\theta_m(\delta_{\omega^1})$ and $a^1 \in A^\theta_m(\delta_{\omega^2})$, because otherwise only one action will be played in the limit.[17] Thus, there exists a mixed belief over ω^1 and ω^2 that makes the myopic agent indifferent between the actions. Since both ω^1 and ω^2 are KL minimizers when $\beta \in [\beta_1, \beta_2]$, any mixed action with $\beta \in [\beta_1, \beta_2]$ is a BN-E, supported by the aforementioned mixed belief. Therefore, the agent's action frequency is almost surely arbitrarily close to the set of Berk-Nash equilibria infinitely often. The argument for other cases is analogous.

[16] This does not hold when $|\mathcal{A}| \geq 3$ because there can be multiple paths connecting any two mixed actions. In fact, Example 2 in Esponda et al. (2019) describes a setting with $|\mathcal{A}| = 3$, in which the dogmatic modeler's action frequency almost surely oscillates around the unique Berk-Nash equilibrium but remains bounded away from it.

[17] δ_{ω^1} and δ_{ω^2} denote the degenerate beliefs at ω^1 and ω^2, respectively.

I conclude this section by returning to Examples 2 and 3.

Example 2, cont. The augmented model θ' is not p'-constrained locally robust. In the unique BN-E, the investor chooses the the high investment level with probability 1, generating a true variance of y higher than the agent's conjecture. In particular, we have

$$K^{p}(2,\omega)=\left(-\frac{1}{2}+\frac{2(\omega_1-2)^2+1}{\omega_2}\right)\sqrt{\frac{\omega_2}{2}}. \tag{14}$$

Hence, the KL minimizer at the BN-E, $(2,1)$, does not locally minimize the KL divergence in the expanded domain Ω^p. In fact, θ does not persist under any prior against the competing model $\theta^c \in \Theta^p$ I previously constructed with $\Omega^{\theta^c}=\{1,2\}\times\{1+\epsilon\}$ when $\epsilon<1$. Theorem 5 then implies that θ' is not p'-constrained locally robust. Interestingly, this example reveals that while a model can be constrained locally robust along one dimension, it may not be so along another dimension.

Example 3, cont. There is a unique self-confirming equilibrium when the supporting belief can take support on any parameters in Ω^p. The equilibrium action profile assigns probability 1 to the pure action $a=2$, and it can be supported by any degenerate belief over the set of KL minimizers, $\Omega^p(\delta_a)\equiv\{\omega\in\Omega:\omega_2\geq 3,\omega_1=2\omega_2+30\}$. This equilibrium is weak if and only if the supporting belief assigns probably 1 to $\tilde{\omega}=(36,3)\in\Omega^p(\delta_a)$. By Theorems 1 and 5, θ is globally robust if $\Omega^\theta\cap(\Omega^p(\delta_a)\setminus\{\tilde{\omega}\})\neq\emptyset$. On the other hand, θ is not p-constrained locally robust if $\Omega^\theta\cap\Omega^p(\delta_a)=\emptyset$, because in this particular environment, θ is locally KL-minimizing at a BN-E σ only if σ is self-confirming. [18] Therefore, the model with the restricted parameter space $\Omega^\theta\subseteq[12,32]\times[1,3]$ is not p-constrained locally robust. Nyarko (1991) and subsequent papers establish that the misspecification described in Example 3 leads to a perpetual cycle between a low price and a high price. However, our analysis reveals that such misspecification is not persistent even when the agent only examines closely-related neighbor models.

5 Extension: Strong Robustness

The various notions of robustness I defined in Section 4 only require that for each particular set of competing models, there exist a prior π_0^θ and a policy a^θ such that θ persists. In

[18] Now suppose we have the borderline case in which $\Omega^\theta\cap\Omega^p(\delta_a)=\{\tilde{\omega}\}$, then whether θ is locally robust depends on other parameters in Ω^θ. For example, if $\Omega^\theta=\{\tilde{\omega}\}$, then $a=2$ is always a best response, thereby satisfying the p-absorbing condition required to apply Theorem 3. However, if $\Omega^\theta=\{\tilde{\omega},(40,2)\}$, then only the high action can be a best response, implying that the action sequence converges to $a=10$; by Theorem 4, θ is not p-constrained locally robust.

$q^*(y^1\|a)$	a^1	a^2
	.5	.5

$q^\theta(y^1\|a,\omega)$	a^1	a^2
ω^1	.5	.7
ω^2	.5	.5

$q^{\theta^c}(y^1\|a,\omega)$	a^1	a^2
ω^3	.5	.5

Table 1: The probability of y^1 given by the true DGP (left), the initial model θ (middle), and the competing model θ^c (right).

this extension, I consider a stronger notion of global robustness that ensures the possibility of long-term adoption regardless of the agent's initial conditions. I then provide a partial characterization that clarifies its content and implications.

Definition 10. A model $\theta \in \Theta$ is *strongly robust* if θ persists in every finite $\Theta^\dagger = \{\theta\} \cup \Theta^c \subset \Theta$ under all full-support $\boldsymbol{\pi}_0^{\Theta^\dagger}$ and optimal $\boldsymbol{a}^{\Theta^\dagger}$.

I start with Example 4, which demonstrates why even a correctly specified θ is not necessarily strongly robust, contrasting Corollary 1.

Example 4 (A globally robust subjective model that fails to be strongly robust)**.** Consider a binary-action-outcome setup where $\mathcal{A} = \{a^1, a^2\}$ and $\mathcal{Y} = \{y^1, y^2\}$. The agent obtains a payoff of 1 if the realized outcome is y_1 and 0 otherwise. The true DGP generates y^1 and y^2 with equal probabilities regardless of the action. The agent entertains a correctly-specified model θ with a binary parameter space $\Omega^\theta = \{\omega^1, \omega^2\}$, with its associated outcome distributions summarized in Table 1. It immediately follows from Corollary 1 that θ is globally robust.

I now construct a competing model against which θ does not persist under some π_0^θ. Let θ^c solely consists of the true DGP, with $\Omega^{\theta^c} = \{\omega^3\}$ and $q^{\theta^c}(\cdot|a,\omega^3) = q^*(\cdot|a), \forall a \in \mathcal{A}$. Besides, let the prior under θ be given by $\pi_0^\theta(\omega^1) = .9$ and take the switching threshold to be $\alpha = 5$. The fact that a^1 is weakly dominated under θ implies that the agent only plays a^2 under regime θ. Consider paths where the agent eventually stops switching and adopts θ. Since ω^2 perfectly matches the outcome distribution but ω^1 does not, the likelihood ratio of ω^1 to ω^2 almost surely converges to 0. Thus, the likelihood ratio $\lambda_t^{\theta^c}$ almost surely converges to $10 > \alpha$ and eventually triggers a switch.[19] It follows that m_t a.s. does not converge to θ. Therefore, θ is not strongly robust.

[19] Formally,

$$\begin{aligned}\lambda_t^{\theta^c} &= \frac{\prod_{\tau=0}^{t-1} q^{\theta^c}(y_\tau|a_\tau,\omega^3)}{\prod_{\tau=0}^{t-1} q^\theta(y_\tau|a_\tau,\omega^1)\,\pi_0^\theta(\omega^1) + \prod_{\tau=0}^{t-1} q^\theta(y_\tau|a_\tau,\omega^2)\,\pi_0^\theta(\omega^2)}\\ &= \frac{1}{\left(\frac{0.7}{0.5}\right)^{N_{t-1}}\left(\frac{0.3}{0.5}\right)^{t-1-N_{t-1}}\pi_0^\theta(\omega^1) + \pi_0^\theta(\omega^2)}\end{aligned}$$

where N_{t-1} is the number of y_1 in the first $t-1$ periods. By the Law of Large Numbers, $\lambda_t^{\theta^c} \to \frac{1}{\pi_0^\theta(\omega^2)} = 10$.

$q^*(y^1\|a)$	a^1	a^2
	.5	.7

$q^\theta(y^1\|a,\omega)$	a^1	a^2
ω^1	.5	.3
ω^2	.5	.4

$q^{\theta^c}(y^1\|a,\omega)$	a^1	a^2
ω^3	.5	.7

Table 2: The probability of y_1 given by the true DGP (left), the subjective model θ (middle), and the competing model θ^c (right).

The key force that prevents θ from being strongly robust is that the Bayes factor rule tends to favor models with a simpler structure. In Example 4, parameter ω^1 is "redundant" in the sense that every supporting belief of every SCE under θ must assign zero probability to this parameter. Therefore, any prior probability assigned to the redundant parameter is a waste and asymptotically bounds the likelihood ratio from above.

It is tempting to conjecture that strong robustness is too strong to tolerate any misspecification. However, this is not true either: a strongly robust model can be both misspecified and inefficient. This claim is substantiated by Example 5 below.

Example 5 (A misspecified yet strongly robust model)**.** Let's use the same binary setup as Example 4 but make a few changes to the conditional probabilities. As shown in Table 2, a^1 is the strictly dominant action under θ. Thus, regardless of the initial conditions, the agent always plays a^1 under θ. Let the true DGP be the competing model. Since a^1 is a self-confirming equilibrium supported by any mixed belief over ω^1 and ω^2, the likelihood ratio is a constant throughout, i.e. $\lambda_t^{\theta^c} = 1, \forall t$. It is obvious that θ also persists against other types of competing models. Therefore, θ is strongly robust and still induces an inefficient action.

Theorem 6 generalizes the intuition we can glean from Examples 4 and 5.

Theorem 6. *The following statements are true:*

(i) If a model $\theta \in \Theta$ is strongly robust, then for every $\omega \in \Omega^\theta$, there must exist an SCE with supporting belief π such that $\omega \in \mathrm{supp}(\pi)$.

(ii) Given any model $\theta \in \Theta$, if $q^\theta(\cdot|a,\omega) \equiv q(\cdot|a)$ for all $a \in \cup_{\pi\in\Delta\Omega^\theta} A_m^\theta(\pi)$ and all $\omega \in \Omega^\theta$, then θ is strongly robust.

Theorem 6(i) provides a necessary condition for strong robustness: no parameter ω under θ is "redundant" in the sense that no SCE supporting belief assigns positive probability to ω. On the other hand, Theorem 6(ii) provides a sufficient condition: if every undominated action is an SCE that can be supported by every belief in $\Delta\Omega^\theta$, then θ is strongly robust.

6 Concluding Remarks

In this paper, I develop and characterize three different robustness criteria of subjective models. Defined based on the chance of long-term persistence against competing models, they provide a direct assessment as to which forms of misspecification are likely to persist.

Instead of assuming that the agent starts outright from a Berk-Nash equilibrium and compares how models fit the data there, this framework incorporates model switching into full-fledged learning dynamics. The characterization highlights the importance of this consideration. For example, global robustness not only needs the existence of a self-confirming equilibrium but also needs it to be p-absorbing. This connects the notion of model robustness with the stability of equilibria under a single model.

The three robustness criteria can be ranked in terms of how hard to satisfy their requirements. This provides a language to compare different models in their degree of robustness. While a model may fail to be globally robust or unconstrained robust, it may be constrained locally robust with respect to a particular parametric family. Varying the size of the family also varies the degree of robustness.

Within our general framework of model switching, there are many other interesting questions to pursue. For example, while global robustness requires a positive chance of a model being adopted forever, we may look at when a model is adopted infinitely often such that the underlying misspecification never vanishes. It may also be interesting to restrict attention to a certain class of models and fully characterize the dynamic patterns of model choices, e.g. how the agent perpetually oscillates between models, and derive more precise predictions about the agent's long-term behavior.

A Auxiliary Results

Lemma 3. *Fix any $\theta, \theta^c \in \Theta$. If θ^c is correctly specified, then as $t \to \infty$, $l_t^\theta / l_t^{\theta^c}$ a.s. converges to ι, where ι is a non-negative random variable with $\mathbb{E}\iota < \infty$.*

Proof. Let $\iota_t = l_t^\theta / l_t^{\theta^c}$, then $\iota_0 = 1, \iota_t \geq 0, \forall t$. I now construct the probability space in which ι_t is a martingale. Given prior $\pi_0^{\theta^c}$, denote by $\mathbb{P}_S^{\theta^c}$ the joint probability measure over Ω^{θ^c} and the set of histories H. In particular, for any $\hat{\Omega} \subset \Omega^{\theta^c}$ and any $\hat{H} \subset H$, we have $\mathbb{P}_S^{\theta^c}(\Omega^p \times H) = \sum_{\omega \in \hat{\Omega}} \pi_0^{\theta^c}(\omega) \mathbb{P}_S^{\theta^c, \omega}(H)$, where $\mathbb{P}_S^{\theta^c, \omega}$ is the probability measure over H induced by the switcher if the true DGP is as described by θ^c and ω. Then,

$$
\begin{aligned}
& \mathbb{E}^{\mathbb{P}_S^{\theta^c}}(\iota_t | h_t) \\
&= \mathbb{E}^{\mathbb{P}_S^{\theta^c}}\left[\frac{\sum_{\omega \in \Omega^\theta} q^\theta(y_{t-1}|a_{t-1}, \omega) \pi_{t-1}^\theta(\omega)}{\sum_{\omega' \in \Omega^{\theta^c}} q^{\theta^c}(y_{t-1}|a_{t-1}, \omega') \pi_{t-1}^{\theta^c}(\omega')} \iota_{t-1} | h_t\right] \\
&= \iota_{t-1} \sum_{\tilde{\omega} \in \Omega^{\theta^c}} \pi_{t-1}^{\theta^c}(\tilde{\omega}) \left[\int_{\mathcal{Y}} \frac{\sum_{\omega \in \Omega^\theta} q^\theta(y_{t-1}|a_{t-1}, \omega) \pi_{t-1}^\theta(\omega)}{\sum_{\omega' \in \Omega^{\theta^c}} q^{\theta^c}(y_{t-1}|a_{t-1}, \omega') \pi_{t-1}^{\theta^c}(\omega')} q^{\theta^c}(y_{t-1}|a_{t-1}, \tilde{\omega}) \nu(dy_{t-1})\right] \\
&= \iota_{t-1} \int_{\mathcal{Y}} \left[\frac{\sum_{\omega \in \Omega^\theta} q^\theta(y_{t-1}|a_{t-1}, \omega) \pi_{t-1}^\theta(\omega)}{\sum_{\omega' \in \Omega^{\theta^c}} q^{\theta^c}(y_{t-1}|a_{t-1}, \omega') \pi_{t-1}^{\theta^c}(\omega')} \left(\sum_{\tilde{\omega} \in \Omega^{\theta^c}} q^{\theta^c}(y_{t-1}|a_{t-1}, \tilde{\omega}) \pi_{t-1}^{\theta^c}(\tilde{\omega})\right)\right] \nu(dy_{t-1}) \\
&= \iota_{t-1} \int_{\mathcal{Y}} \left[\sum_{\omega \in \Omega^\theta} q^\theta(y_{t-1}|a_{t-1}, \omega) \pi_{t-1}^\theta(\omega)\right] \nu(dy_{t-1}) \\
&= \iota_{t-1} \sum_{\omega \in \Omega^\theta} \left[\int_{\mathcal{Y}} q^\theta(y_{t-1}|a_{t-1}, \omega) \nu(dy_{t-1})\right] \pi_{t-1}^\theta(\omega) d\omega = \iota_{t-1}.
\end{aligned}
$$

Hence, ι_t is a martingale w.r.t. $\mathbb{P}_S^{\theta^c}$. Since $\iota_t \geq 0, \forall t$, the Martingale Convergence Theorem implies that ι_t converges to ι almost surely w.r.t. $\mathbb{P}_S^{\theta^c}$, and $\mathbb{E}^{\mathbb{P}_S^{\theta^c}} \iota \leq \mathbb{E}^{\mathbb{P}_S^{\theta^c}} \iota_0 = 1$. Since θ^c is correctly specified, there exists a parameter $\omega^* \in \Omega^{\theta^c}$ such that $q^*(\cdot|a) \equiv q^{\theta^c}(\cdot|a, \omega^*), \forall a \in \mathcal{A}$. It then follows from $\pi_0^{\theta^c}(\omega^*) > 0$ that ι_t also converges to ι almost surely w.r.t. $\mathbb{P}_S^{\theta^c, \omega^*} \equiv \mathbb{P}_S$. Moreover, $\mathbb{E}\iota < \infty$ because otherwise it contradicts $\mathbb{E}^{\mathbb{P}_S^{\theta^c}} \iota \leq 1$. $\square$

Lemma 4. *Fix any $\theta, \theta' \in \Theta$, $\omega \in \Omega^\theta, \omega' \in \Omega^{\theta'}$ and any sequence of actions $(a_1, a_2, ...)$. For each infinite history $h \in (\mathcal{A} \times \mathcal{Y})^\infty$ that is generated according to $(a_1, a_2, ...)$ by the true DGP, let*

$$
\xi_t(h) = \ln \frac{q^\theta(y_t|a_t, \omega)}{q^{\theta'}(y_t|a_t, \omega')} - \mathbb{E}\left(\ln \frac{q^\theta(y_t|a_t, \omega)}{q^{\theta'}(y_t|a_t, \omega')} | h_t\right).
$$

Then for any fixed $t_0 \geq 1$,

$$
\lim_{t \to \infty} (t - t_0 + 1)^{-1} \sum_{\tau = t_0}^{t} \xi_\tau(h) = 0, \text{ a.s..}
$$

Proof. $\xi_t(h)$ is a martingale difference process since $E(\xi_t(h)|h_t) = 0$. Hence, for any t_0, $\xi_{t_0}^t(h) \coloneqq \sum_{\tau=t_0}^t (t-\tau+1)^{-1}\xi_\tau(h)$ is also a martingale difference process. To use the Martingale Convergence Theorem, I now show that $\sup_t \mathbb{E}\left(\left(\xi_{t_0}^t\right)^2\right) < \infty$. Notice that

$$\begin{aligned}
\mathbb{E}\left(\left(\xi_{t_0}^t\right)^2\right) =& \mathbb{E}\left[\left(\sum_{\tau=t_0}^t (t-\tau+1)^{-1}\xi_\tau(h)\right)^2\right] \\
\leq& \sum_{\tau=t_0}^t (t-\tau+1)^{-2}\,\mathbb{E}\left[(\xi_\tau(h))^2\right] \\
\leq& \sum_{\tau=t_0}^t (t-\tau+1)^{-2}\,\mathbb{E}\left[\left(\ln\frac{q^\theta(y_t|a_t,\omega)}{q^{\theta'}(y_t|a_t,\omega')}\right)^2\right] \\
\leq& \sum_{\tau=t_0}^t (t-\tau+1)^{-2}\,\mathbb{E}\left[\left(\ln\frac{q^*(y_t|a_t)}{q^\theta(y_t|a_t,\omega)}\right)^2 + \left(\ln\frac{q^*(y_t|a_t)}{q^{\theta'}(y_t|a_t,\omega')}\right)^2\right] \\
\leq& 2\sum_{\tau=t_0}^t (t-\tau+1)^{-2}\max_a \mathbb{E}\left[(g_a(y))^2\right] < \infty,
\end{aligned}$$

where the first inequality follows from the fact that, for any $\tau' > \tau \geq t_0$, $\mathbb{E}(\xi_\tau(h)\xi_{\tau'}(h)) = \mathbb{E}(\mathbb{E}(\xi_{\tau'}(h)|h_{\tau'})\xi_\tau(h)) = 0$ and the last inequality follows from Assumption 2. Now we can invoke the Martingale Convergence Theorem which implies that $\xi_{t_0}^t$ converges to a random variable $\xi_{t_0}^\infty$ almost surely with $\mathbb{E}\left(\left(\xi_{t_0}^\infty\right)^2\right) < \infty$. By Kronecker Lemma, since $\xi_{t_0}^\infty = \lim_{t\to\infty}\sum_{\tau=t_0}^t (t-\tau+1)^{-1}\xi_\tau(h)$ is finite a.s., we have

$$\lim_{t\to\infty}(t-t_0+1)^{-1}\sum_{\tau=t_0}^t \xi_\tau(h) = 0, \text{a.s..}$$

□

The action frequency $\sigma_t : \mathcal{A}^t \to \Delta\mathcal{A}$ measures how frequent each action has been played up to period t. In particular, given an action sequence $(a_0, a_1, ...)$,

$$\sigma_t(a) = \frac{\sum_{\tau=0}^{t-1}\mathbf{1}(a_t = a)}{t}.$$

Lemma 5. *Suppose the action frequency of a dogmatic modeler with $\Theta^\dagger = \{\theta\}$ converges to σ, then the dogmatic modeler's belief $\tilde{\pi}_t^\theta$ a.s. converges to $\tilde{\pi}^\theta$, with $\tilde{\pi}^\theta\left(\Omega^\theta(\sigma)\right)$ a.s. converges to 1. Similarly, if the action frequency of a switcher with $\Theta^\dagger \ni \theta$ converges to σ, then her belief π_t^θ also a.s. converges to some π^θ with $\pi^\theta\left(\Omega^\theta(\sigma)\right)$.*

Proof. I now show that the claim is true for a switcher; the proof for a dogmatic modeler is completely identical. Since Ω^θ is finite, for any given σ, there exists $\epsilon > 0$ such that

$$\sum_{\mathcal{A}} \sigma(a) \left[D_{KL}\left(q^*(\cdot|a) \parallel q^\theta(\cdot|a,\omega)\right) - D_{KL}\left(q^*(\cdot|a) \parallel q^\theta(\cdot|a,\omega')\right)\right] < -\epsilon, \tag{15}$$

for all $\omega \in \Omega^\theta(\sigma)$ and $\omega' \in \Omega^\theta / \Omega^\theta(\sigma)$. Consider the agent's belief over any $\omega \in \Omega^\theta(\sigma)$ and $\omega' \in \Omega^\theta / \Omega^\theta(\sigma)$ at time t,

$$\begin{aligned} \frac{\pi_t^\theta(\omega')}{\pi_t^\theta(\omega)} &= \frac{\prod_{\tau=0}^{t-1} q^\theta(y_\tau|a_\tau,\omega')\,\pi_0^\theta(\omega')}{\prod_{\tau=0}^{t-1} q^\theta(y_\tau|a_\tau,\omega)\,\pi_0^\theta(\omega)} \\ &= \exp\left(\sum_{\tau=0}^{t-1} \ln \frac{q^\theta(y_\tau|a_\tau,\omega')}{q^\theta(y_\tau|a_\tau,\omega)} + \ln \frac{\pi_0^\theta(\omega')}{\pi_0^\theta(\omega)}\right). \end{aligned}$$

We are done if this ratio converges to 0. Notice that

$$\begin{aligned} &\frac{1}{t}\sum_{\tau=0}^{t-1} \mathbb{E}\left(\ln \frac{q^\theta(y_\tau|a_\tau,\omega')}{q^\theta(y_\tau|a_\tau,\omega)} | h_t\right) \\ =& -\sum_{\mathcal{A}} \sigma_t(a)\left[D_{KL}\left(q^*(\cdot|a) \parallel q^\theta(\cdot|a,\omega')\right) - D_{KL}\left(q^*(\cdot|a) \parallel q^\theta(\cdot|a,\omega)\right)\right], \end{aligned}$$

which converges to the left-hand side of Eq. (15) as σ_t converges to σ. Hence, there exists T_1 such that

$$\frac{1}{t}\sum_{\tau=0}^{t-1} \mathbb{E}\left(\ln \frac{q^\theta(y_\tau|a_\tau,\omega')}{q^\theta(y_\tau|a_\tau,\omega)} | h_t\right) < -\frac{\epsilon}{2}, \forall t > T_1.$$

By Lemma 4, there exists T_2 such that when $t > T_2$,

$$\frac{1}{t}\sum_{\tau=0}^{t-1} \ln \frac{q^\theta(y_\tau|a_\tau,\omega')}{q^\theta(y_\tau|a_\tau,\omega)} < \frac{1}{t}\sum_{\tau=0}^{t-1} \mathbb{E}\left(\ln \frac{q^\theta(y_\tau|a_\tau,\omega')}{q^\theta(y_\tau|a_\tau,\omega)} | h_t\right) + \frac{\epsilon}{3}$$

It follows that when $t > \max\{T_1, T_2\}$,

$$\sum_{\tau=0}^{t-1} \ln \frac{q^\theta(y_\tau|a_\tau,\omega')}{q^\theta(y_\tau|a_\tau,\omega)} < t \cdot \left(-\frac{\epsilon}{6}\right).$$

Hence, $\frac{\pi_t^\theta(\omega')}{\pi_t^\theta(\omega)}$ converges to 0. □

Lemma 6. *Fix any $\theta \in \Theta$, the optimal action correspondence $A^\theta : \Delta\Omega^\theta \rightrightarrows \mathcal{A}$ is upper hemicontinuous in both the belief π and the discount factor δ.*

Proof. This is a standard result directly following from Blackwell (1965) and Maitra (1968).

□

Lemma 7. *Fix any $\theta \in \Theta$, the set of all Berk-Nash equilibria under θ is compact.*

Proof. Denote the set of all Berk-Nash equilibria under model θ as $BN^{\theta} \subset \Delta\mathcal{A}$. Since $\Delta\mathcal{A}$ is bounded, we only need to show that BN^{θ} is closed. Suppose σ is the limit of some sequence $(\sigma_n)_n$ of Berk-Nash equilibria, but σ is not a Berk-Nash equilibrium, i.e. $\sigma \notin BN^{\theta}$. Then for every belief $\pi \in \Delta\Omega^{\theta}(\sigma)$, we have that $\sigma \notin \Delta A_m^{\theta}(\pi)$. Since $\Omega^{\theta}(\cdot)$ is upper hemicontinuous, it must be that $\Omega^{\theta}(\sigma_n) \subset \Omega^{\theta}(\sigma)$ for large enough n. Hence, we have $\sigma \notin \Delta A_m^{\theta}(\pi)$ for every belief $\pi \in \Delta\Omega^{\theta}(\sigma_n)$ when n is large enough. However, we know that $\text{supp}(\sigma) \subset \text{supp}(\sigma_n)$ for large enough n, which implies that $\sigma_n \notin \Delta A_m^{\theta}(\pi)$ for large n. This is a contradiction. □

B Main Results

B.1 Proof of Theorem 1

I first prove Lemmas 1 and 2.

Proof of Lemma 1. That $l_t^{\theta^c}/l_t^{\theta}$ a.s. converges to $\iota \leq \alpha$ on paths where m_t converges to θ immediately follows from Lemma 3. I now show that π_t^{θ} and $\pi_t^{\theta^c}$ also a.s. converge. Given any $\omega \in \Omega^{\theta}$, we can decompose $\pi_t^{\theta}(\omega)$ as follows,

$$\begin{aligned}\frac{\pi_t^{\theta}(\omega)}{\pi_0^{\theta}(\omega)} &= \frac{\prod_{\tau=0}^{t-1} q^{\theta}(y_\tau|a_\tau,\omega)}{\sum_{\omega'\in\Omega^{\theta}}\prod_{\tau=0}^{t-1} q^{\theta}(y_\tau|a_\tau,\omega')\,\pi_0^{\theta}(\omega')} \\ &= \frac{l_t^{\theta^c}}{l_t^{\theta}}\cdot\frac{\prod_{\tau=0}^{t-1} q^{\theta}(y_\tau|a_\tau,\omega)}{\sum_{\omega''\in\Omega^{\theta^c}}\prod_{\tau=0}^{t-1} q^{\theta^c}(y_\tau|a_\tau,\omega'')\,\pi_0^{\theta^c}(\omega'')} \\ &:= \frac{l_t^{\theta^c}}{l_t^{\theta}}\cdot\frac{l_t^{\theta,\omega}}{l_t^{\theta^c}},\end{aligned}$$

where the second term $l_t^{\theta,\omega}/l_t^{\theta^c}$ is the likelihood ratio of the competing model θ^c and a model that consists of a single parameter ω. By Lemma 3, $l_t^{\theta,\omega}/l_t^{\theta^c}$ a.s. converges. Consider the paths on which m_t converges to θ. On these paths, both $l_t^{\theta^c}/l_t^{\theta}$ and $l_t^{\theta,\omega}/l_t^{\theta^c}$ converges a.s., which implies that $\pi_t^{\theta}(\omega)$ a.s. converges as well. Since this is true for all $\omega \in \Omega^{\theta}$, π_t^{θ} a.s. converges to some limit π_∞^{θ} on those paths. Analogously, for any $\omega' \in \Omega^{\theta^c}$, we can decompose $\pi_t^{\theta^c}(\omega')$ as follows,

$$\frac{\pi_t^{\theta^c}(\omega')}{\pi_0^{\theta^c}(\omega')} = \frac{\prod_{\tau=0}^{t-1} q^{\theta^c}(y_\tau|a_\tau,\omega')}{\sum_{\omega''\in\Omega^{\theta^c}}\prod_{\tau=0}^{t-1} q^{\theta^c}(y_\tau|a_\tau,\omega'')\,\pi_0^{\theta^c}(\omega'')},$$

which, again by Lemma 3, converges almost surely. □

Proof of Lemma 2. Consider any $\hat{\omega}$ such that with positive probability, m_t converges to θ and $\hat{\omega} \in \text{supp}\left(\pi_\infty^\theta\right)$. Let $A^-(\hat{\omega}) \equiv \left\{a \in \mathcal{A} : q^\theta(\cdot|a,\hat{\omega}) \neq q^*(\cdot|a)\right\}$. I now show that every action in $A^-(\hat{\omega})$ is played at most finite times a.s. on the paths where m_t converges to θ. Suppose instead that actions in $A^-(\hat{\omega})$ are played infinitely often. Then there must exist some $\gamma > 0$ such that $\mathbb{E}\ln\frac{q^*(y|a_t)}{q^\theta(y|a_t,\hat{\omega})} > \gamma$ for infinitely many t. Since θ^c is correctly specified, there exists a parameter $\omega^* \in \Omega^{\theta^c}$ such that $q^*(\cdot|a) \equiv q^{\theta^c}(\cdot|a,\omega^*), \forall a \in \mathcal{A}$. Hence, $\mathbb{E}\ln\frac{q^{\theta^c}(y|a_t,\omega^*)}{q^\theta(y|a_t,\hat{\omega})} > \gamma$ for infinitely many t. Notice that

$$\begin{aligned}
\frac{l_t^{\theta^c}}{l_t^\theta} &= \frac{\sum_{\omega' \in \Omega^{\theta^c}} \prod_{\tau=0}^{t-1} q^{\theta^c}(y_\tau|a_\tau,\omega')\,\pi_0^{\theta^c}(\omega')}{\sum_{\omega \in \Omega^\theta} \prod_{\tau=0}^{t-1} q^\theta(y_\tau|a_\tau,\omega)\,\pi_0^\theta(\omega)} \\
&> \pi_t^\theta(\hat{\omega})\frac{\pi_0^{\theta^c}(\omega^*)}{\pi_0^\theta(\hat{\omega})}\frac{\prod_{\tau=0}^{t-1} q^{\theta^c}(y_\tau|a_t,\omega^*)}{\prod_{\tau=0}^{t-1} q^\theta(y_\tau|a_\tau,\hat{\omega})} \\
&> \pi_t^\theta(\hat{\omega})\frac{\pi_0^{\theta^c}(\omega^*)}{\pi_0^\theta(\hat{\omega})}\left[\sum_{\tau=0}^{t-1} 1_{\{a_\tau \in A^-(\hat{\omega})\}}\ln\frac{q^{\theta^c}(y_\tau|a_t,\omega^*)}{q^\theta(y_\tau|a_\tau,\hat{\omega})}\right],
\end{aligned}$$

which almost surely increase to infinity by Lemma 4 as $t \to \infty$, contradicting the assumption that m_t converges to θ. Therefore, on the paths where $m_t \to \theta$, almost surely, there exists T such that $a_t \in \mathcal{A}\setminus \cup_{\omega' \in \text{supp}(\pi_\infty^\theta)} A^-(\hat{\omega}), \forall t > T$.

Since $q^\theta(\cdot|a,\omega') \equiv q^*(\cdot|a)$ for all $\omega' \in \text{supp}\left(\pi_\infty^\theta\right)$ and all $a \in \mathcal{A}\setminus \cup_{\omega' \in \text{supp}(\pi_\infty^\theta)} A^-(\omega')$, the actions that are played in the limit have no experimentation value and are myopically optimal. Therefore, fix a particular value of π_∞^θ that is a limit belief for a positive measure of histories where $m_t \to \theta$, there exists a set of actions $A \subset A_m^\theta\left(\pi_\infty^\theta\right)$ such that with positive probability, the agent only plays actions from this set in the limit. Since m_t eventually converges to θ, it must be true that with positive probability, a dogmatic modeler also only plays actions from A in the limit. Therefore, any strategy σ with $\text{supp}(\sigma) = A$ is a p-absorbing self-confirming equilibrium under θ. □

Suppose there exists a p-absorbing SCE σ under θ. Consider the learning process of a dogmatic modeler with $\Theta^\dagger = \{\theta\}$. There exists a full-support prior $\pi_0^\theta \in \Delta\Omega^\theta$ and an optimal policy a^θ such that with positive probability, she eventually only chooses actions from $\text{supp}(\sigma)$ and each element of $\text{supp}(\sigma)$ is played infinitely often (if there exists $a \in \text{supp}(\sigma)$ s.t. a is only played finite times, then we can find an SCE σ' with a smaller support such that each element of $\text{supp}(\sigma')$ is played i.o.). Denote those paths by $\tilde{H}$. Then by a similar argument as in the proof of Lemma 2, π_t^θ a.s. converges to a limit π_∞^θ on $\tilde{H}$, with $\text{supp}\left(\pi_\infty^\theta\right) \subseteq G_\sigma^\theta := \{\omega \in \Omega^\theta : q^*(\cdot|a) = q^\theta(\cdot|a,\omega), \forall a \in \text{supp}(\sigma)\}$.

Now take any finite $\Theta^c \subset \Theta$ and any prior $\boldsymbol{\pi}_0^{\Theta^c}$. For every $\theta' \in \Theta^c$, denote the set of parameters in $\Omega^{\theta'}$ that yield a zero KL divergence when σ is played as $G_\sigma^{\theta'}$. That is,

$G_\sigma^{\theta'} \coloneqq \left\{\omega \in \Omega^{\theta'} : q^*(\cdot|a) = q^{\theta'}(\cdot|a,\omega), \forall a \in \operatorname{supp}(\sigma)\right\}$. Note that this set may be empty. Define $\eta_{T,t}(\theta',\omega') \coloneqq \prod_{\tau=T}^{t} \frac{q^{\theta'}(y_\tau|a_\tau,\omega')}{q^*(y_\tau|a_\tau)}$. Let $H^\sigma \coloneqq \{h \in H : a_t \in \operatorname{supp}(\sigma), \forall t \geq 0\}$ denote the set of histories where all actions are taken from the support of σ. Notice that $\hat{H} \subseteq H^\sigma$. For all $\omega' \in \Omega^{\theta'} \backslash G_\sigma^{\theta'}$, there exists some $\gamma > 0$ such that either $q^{\theta'}(y|a,\omega') \equiv q^*(y|a)$ or $\mathbb{E}\left(\frac{q^{\theta'}(y|a,\omega')}{q^*(y|a)}\right) < -\gamma$ for all $a \in \operatorname{supp}(\sigma)$. By Lemma 4, we have that $\eta_{T,t}(\theta',\omega')$ a.s. converges to 0 on H^σ and thus on $\tilde{H}$.

Therefore, on $\tilde{H}$, the dogmatic modeler eventually only chooses actions from $\operatorname{supp}(\sigma)$, with $\pi_t^\theta \overset{a.s.}{\to} \pi_\infty^\theta$ and $\eta_{T,t}(\theta',\omega') \overset{a.s.}{\to} 0$ for every $\theta' \in \Theta^c$ and $\omega' \in G_\sigma^{\theta'}$. This implies the existence of an integer $T > 0$ such that, with positive probability, we have (1) $a_t \in \operatorname{supp}(\sigma), \forall t \geq T$, (2) π_t^θ converges to a limit π_∞^θ with $\operatorname{supp}\left(\pi_\infty^\theta\right) \subset G_\sigma^\theta$, and (3) $\eta_{T,t}(\theta',\omega') \leq 1$ for all $t \geq T$ and all $\theta' \in \Theta^c$ and $\omega' \in G_\sigma^{\theta'}$. Let $\epsilon > 0$ be small enough such that $\frac{1}{1-\epsilon} < \alpha$. We can find a new full-support prior $\hat{\pi}_0^\theta \in B_\epsilon\left(\Delta G_\sigma^\theta\right)$ under which, on a positive measure of histories $\hat{H}$, a dogmatic modeler sees that (1') $a_t \in \operatorname{supp}(\sigma), \forall t \geq 0$, (2') the posterior $\hat{\pi}_t^\theta$ a.s. converges to $\hat{\pi}_\infty^\theta$ and never leaves $B_\epsilon\left(\Delta G_\sigma^\theta\right), \forall t \geq 0$, and (3') $\eta_{0,t}(\theta',\omega') \leq 1, \forall t \geq 0$ for all $\theta' \in \Theta^c$ and $\omega' \in G_\sigma^{\theta'}$.

Consider any $\theta' \in \Theta^c$ with $\pi_0^{\theta'} \in \Omega^{\theta'}$ and the likelihood ratio $l_t^{\theta'}/l_t^\theta$ computed from any history in $\hat{H}$ under prior $\hat{\pi}_0^\theta$,

$$
\begin{aligned}
\frac{l_t^{\theta'}}{l_t^\theta} &= \frac{\sum_{\omega' \in \Omega^{\theta'}} \pi_0^{\theta'}(\omega') \prod_{\tau=0}^{t-1} q^{\theta'}(y_\tau|a_\tau,\omega')}{\sum_{\omega \in \Omega^\theta} \hat{\pi}_0^\theta(\omega) \prod_{\tau=0}^{t-1} q^\theta(y_\tau|a_\tau,\omega)} \\
&< \frac{\sum_{\omega' \in G_\sigma^{\theta'}} \pi_0^{\theta'}(\omega') \prod_{\tau=0}^{t-1} q^*(y_\tau|a_\tau) + \sum_{\omega'' \in \Omega^{\theta'} \backslash G_\sigma^{\theta'}} \pi_0^{\theta'}(\omega'') \prod_{\tau=0}^{t-1} q^{\theta'}(y_\tau|a_\tau,\omega'')}{\sum_{\omega \in G_\sigma^\theta} \hat{\pi}_0^\theta(\omega) \prod_{\tau=0}^{t-1} q^*(y_\tau|a_\tau)} \\
&= \frac{\pi_0^{\theta'}\left(G_\sigma^{\theta'}\right)}{\hat{\pi}_0^\theta\left(G_\sigma^\theta\right)} + \sum_{\omega'' \in \Omega^{\theta'} \backslash G_\sigma^{\theta'}} \frac{\pi_0^{\theta'}(\omega'')}{\hat{\pi}_0^\theta\left(G_\sigma^\theta\right)} \eta_{0,t}(\theta',\omega'') \\
&\leq \frac{1}{\hat{\pi}_0^\theta\left(G_\sigma^\theta\right)} + \frac{\pi_0^{\theta'}\left(\Omega^{\theta'} \backslash G_\sigma^{\theta'}\right)}{\hat{\pi}_0^\theta\left(G_\sigma^\theta\right)} \\
&\leq \frac{1}{1-\epsilon} < \alpha
\end{aligned}
$$

where the first inequality follows from the definition of G_σ^θ and $G_\sigma^{\theta'}$ and that $\hat{\pi}_0^\theta$ is full-support, the second inequality follows from (3'), and the third inequality from (2'). If for all $\theta' \in \Theta^c$, the likelihood ratio $l_t^{\theta'}/l_t^\theta$ never exceeds the threshold α, then the Bayes factor λ_t^θ calculated from such a history never exceeds α as well. Thus, on any history $h \in \hat{H}$, the switcher never makes any switch to any model $\theta' \in \Theta^c$, i.e. $m_t = \theta, \forall t \geq 0$. Therefore, if we endow the switcher with the same prior $\hat{\pi}_0^\theta$, then $\hat{H}$ also has a positive measure under $\mathbb{P}_S$.

B.2 Proof of Corollary 2

It suffices to show that every uniformly quasi-strict SCE σ is p-absorbing. By definition, there exists a belief $\pi \in \Delta\Omega^{\theta}$ with $\operatorname{supp}(\pi) \subset G_{\sigma}^{\theta}$ (the set G_{σ}^{θ} is defined in the proof of Theorem 1). Since σ is uniformly quasi-strict, $\operatorname{supp}(\sigma)$ contains all myopically optimal actions against each degenerate belief δ_{ω} concentrated on $\omega \in \operatorname{supp}(\pi)$. In addition, $\operatorname{supp}(\sigma)$ must be optimal against δ_{ω} for an agent who maximizes discounted utility, because the dynamic programming problem described by (4) reduces to a static maximization problem when the belief is degenerate. This implies that $\operatorname{supp}(\sigma)$ is also optimal against π. Further, since A^{θ} is upper hemicontinuous (by Lemma 6), there exists $\gamma > 0$ small enough such that $\operatorname{supp}(\sigma) = A^{\theta}(\tilde{\pi})$ for all $\tilde{\pi} \in B_{\gamma}(\pi)$.

Suppose $a_t \in \operatorname{supp}(\sigma), \forall t \geq 0$, then for every $\omega \in \Omega^{\theta} \backslash G_{\sigma}^{\theta}$,

$$\begin{aligned}
\mathbb{E}\left[\frac{\pi_t^{\theta}(\omega)}{\pi_t^{\theta}(G_{\sigma}^{\theta})}|h_t\right] &= \mathbb{E}\left[\frac{\pi_0^{\theta}(\omega)\prod_{\tau=0}^{t-1} q^{\theta}(y_{\tau}|a_{\tau},\omega)}{\sum_{\omega' \in G_{\sigma}^{\theta}} \pi_0^{\theta}(\omega')\prod_{\tau=0}^{t-1} q^{\theta}(y_{\tau}|a_{\tau},\omega')}|h_t\right] \\
&= \mathbb{E}\left[\frac{\pi_0^{\theta}(\omega)}{\pi_0^{\theta}(G_{\sigma}^{\theta})}\frac{\prod_{\tau=0}^{t-1} q^{\theta}(y_{\tau}|a_{\tau},\omega)}{\prod_{\tau=0}^{t-1} q^{*}(y_{\tau}|a_{\tau})}|h_t\right] \\
&= \frac{\pi_0^{\theta}(\omega)\prod_{\tau=0}^{t-2} q^{\theta}(y_{\tau}|a_{\tau},\omega)}{\pi_0^{\theta}(G_{\sigma}^{\theta})\prod_{\tau=0}^{t-2} q^{*}(y_{\tau}|a_{\tau})} = \frac{\pi_{t-1}^{\theta}(\omega)}{\pi_{t-1}^{\theta}(G_{\sigma}^{\theta})}
\end{aligned}$$

Therefore, $\frac{\pi_t^{\theta}(\omega)}{\pi_t^{\theta}\left(G_{\sigma}^{\theta}\right)}$ is a non-negative supermartingale for every $\omega \in \Omega^{\theta} \backslash G_{\sigma}^{\theta}$. It follows that $\frac{\pi_t^{\theta}\left(\Omega^{\theta} \backslash G_{\sigma}^{\theta}\right)}{\pi_t^{\theta}\left(G_{\sigma}^{\theta}\right)}$ is also non-negative supermartingale. By the maximal inequality, for all $\epsilon > 0$,

$$\mathbb{P}_B\left(\frac{\pi_t^{\theta}\left(\Omega^{\theta} \backslash G_{\sigma}^{\theta}\right)}{\pi_t^{\theta}\left(G_{\sigma}^{\theta}\right)} \geq \frac{\pi_0^{\theta}\left(\Omega^{\theta} \backslash G_{\sigma}^{\theta}\right)}{\pi_0^{\theta}\left(G_{\sigma}^{\theta}\right)} + \epsilon \text{ for some } t\right) < 1.$$

Since $\pi_t^{\theta}\left(G_{\sigma}^{\theta}\right) = 1 - \pi_t^{\theta}\left(\Omega^{\theta} \backslash G_{\sigma}^{\theta}\right)$, the above inequality implies that for all $\epsilon > 0$,

$$\mathbb{P}_B\left(\pi_t^{\theta}\left(\Omega^{\theta} \backslash G_{\sigma}^{\theta}\right) \geq \pi_0^{\theta}\left(\Omega^{\theta} \backslash G_{\sigma}^{\theta}\right) + \epsilon \text{ for some } t\right) < 1.$$

Pick some $\gamma' \in (0, \gamma)$ and $\pi_0^{\theta} \in B_{\gamma'}(\pi)$, then $\pi_0^{\theta}\left(G_{\sigma}^{\theta}\right) > 1 - \gamma'$. Notice that the belief ratio $\frac{\pi_t^{\theta}(\omega)}{\pi_t^{\theta}(\omega')}$ remain unchanged throughout all periods provided that $\omega, \omega' \in G_{\sigma}^{\theta}$. Hence, if $\pi_t^{\theta} \notin B_{\gamma}(\pi)$ for some $t \geq 0$, then there exists t such that $\pi_t^{\theta}\left(\Omega^{\theta} \backslash G_{\sigma}^{\theta}\right) \geq \pi_0^{\theta}\left(\Omega^{\theta} \backslash G_{\sigma}^{\theta}\right) + \gamma - \gamma'$.

Therefore,

$$\begin{aligned}&\mathbb{P}_B\left(\pi_t^\theta \notin B_\gamma(\pi) \text{ for some } t \geq 0\right)\\ \leq &\mathbb{P}_B\left(\pi_t^\theta\left(\Omega^\theta \backslash G_\sigma^\theta\right) \geq \pi_0^\theta\left(\Omega^\theta \backslash G_\sigma^\theta\right)+\gamma-\gamma' \text{ for some } t\right)<1.\end{aligned}$$

This implies that $\mathbb{P}_B\left(\pi_t^\theta \in B_\gamma(\pi), \forall t \geq 0\right)>0$. Notice that $\pi_t^\theta \in B_\gamma(\pi), \forall t \geq 0$ in turn implies that a dogmatic modeler will only play actions from $\operatorname{supp}(\sigma)$. Therefore, σ is p-absorbing.

B.3 Proof of Theorem 2

I now show that if θ is unconstrained locally robust, then it must persist against a correctly specified model under some priors and policies. From there, we can use Lemma 2 and Theorem 1 to show the equivalence between unconstrained local robustness and global robustness.

Denote the parameter set of θ as $\Omega^\theta=\left\{\omega^1, \ldots, \omega^N\right\}$. Consider a competing model θ^c constructed as below:

- $\Omega^{\theta^c}=\Omega^\theta$
- $q^{\theta^c}\left(\cdot \mid a, \omega^n\right)=(1-\epsilon) q^\theta\left(\cdot \mid a, \omega^n\right)+\epsilon q^*(\cdot \mid a), \forall a \in \mathcal{A}, \forall \omega^n \in \Omega^\theta$

By construction, $\theta^c \in N_\epsilon(\theta)$. Hence, there exists $\epsilon>0$ such that θ persists against θ^c under some full-support priors π_0^θ, $\pi_0^{\theta^c}$ and some policies a^θ, a^{θ^c}. Further, this implies that there exists $\epsilon>0$ and some initial condition such that the probability that $l_t^{\theta^c} / l_t^\theta \leq \alpha$ for all $t \geq 0$ is strictly positive. Observe that

$$\begin{aligned}\frac{l_t^{\theta^c}}{l_t^\theta} & =\frac{\sum_{\omega \in \Omega^{\theta^c}} \pi_0^{\theta^c}(\omega) \prod_{\tau=0}^{t-1} q^{\theta^c}\left(y_\tau \mid a_\tau, \omega\right)}{\sum_{\omega \in \Omega^\theta} \pi_0^\theta(\omega) \prod_{\tau=0}^{t-1} q^\theta\left(y_\tau \mid a_\tau, \omega\right)} \\ & =1-\epsilon+\epsilon \frac{\prod_{\tau=0}^{t-1} q^*\left(y_\tau \mid a_\tau\right)}{\sum_{\omega \in \Omega^\theta} \pi_0^\theta(\omega) \prod_{\tau=0}^{t-1} q^\theta\left(y_\tau \mid a_\tau, \omega\right)} .\end{aligned}$$

Notice that the last term, denoted by l_t^* / l_t^θ, is the likelihood ratio of the true DGP and θ. If there exists $T>0$ such that with positive probability, $l_t^{\theta^c} / l_t^\theta \leq \alpha$ for all $t \geq T$, than it must be that $l_t^* / l_t^\theta \leq \frac{1}{\epsilon}(\alpha+1-\epsilon)$ for all $t \geq T$. Hence, θ persists against the true DGP under switching threshold $\alpha'=\frac{1}{\epsilon}(\alpha+1-\epsilon)>1$. Since Lemma 2 and Theorem 1 hold regardless of the switching threshold, this implies that θ is globally robust.

B.4 Proof of Theorem 3

Suppose σ is a pure p-absorbing BN-E with θ being locally KL-minimizing and locally identified at σ w.r.t. Ω^p, and σ assigns probability 1 to $a^* \in \mathcal{A}$. Then there exists a full-support prior π_0^θ and a policy a^θ such that with this prior, a dogmatic modeler eventually only plays a^* with positive probability. It follows from Lemma (5) that $\pi_t^\theta\left(\Omega^\theta\left(\sigma\right)\right) \overset{a.s.}{\to} 1$. Therefore, for any $\gamma \in (0, 1/\alpha)$, there is a new full-support $\tilde{\pi}_0^\theta$ and a positive measure of paths, denoted by $\tilde{H} \subset H$, where $a_t \in \operatorname{supp}\left(\sigma\right)$ and $\tilde{\pi}_t^\theta\left(\Omega^\theta\left(\sigma\right)\right) > \gamma, \forall t = 0, 1, ...$

For convenience, I write $K^p\left(\sigma, \omega\right)$ as $K^p\left(a^*, \omega\right)$. Pick $\eta > 0$ such that θ satisfies the conditions of being locally KL-minimizing and locally identified at σ. That is, for all $\omega \in \Omega^\theta\left(\sigma\right)$, $\omega' \in B_\eta\left(\omega\right) \cup \Omega^p$, and $a \in \operatorname{supp}\left(\sigma\right)$, we have that either

$$K^p\left(a^*, \omega\right) < K^p\left(a^*, \omega'\right),$$

or

$$p\left(\cdot|a^*, \omega\right) \equiv p\left(\cdot|a^*, \omega'\right).$$

Let $\bar{\epsilon} \in \left(0, \min\left\{\eta, \frac{1}{2}\min_{\omega, \omega' \in \Omega^\theta}\|\omega - \omega'\|\right\}\right)$ be such that for all $\epsilon \leq \bar{\epsilon}$, $\omega \in \Omega^\theta \backslash \Omega^\theta\left(\sigma\right)$ and $\omega' \in B_\epsilon\left(\omega\right) \cap \Omega^p$, we have

$$K^p\left(a^*, \omega'\right) > K^p\left(a^*, \omega''\right), \forall \omega'' \in \Omega^\theta\left(\sigma\right).$$

The existence of such $\bar{\epsilon}$ is guaranteed by the finiteness of Ω^θ and the continuity of $K^p\left(a, \omega\right)$ in ω. This condition requires that if some ω is not a KL minimizer under θ at a, then slightly perturbing ω still yields a strictly higher KL divergence than a KL minimizer does. Let $\epsilon = \bar{\epsilon}$ and fix a set of competing models $\Theta^c \subset \Theta_\epsilon^p\left(\theta\right)$. To ease notation, denote the likelihood of $\omega \in \Omega^p$ as

$$l_t^{p,\omega} = \prod_{\tau=0}^{t-1} p\left(y_\tau | a_\tau, \omega\right).$$

By the definition of ϵ, for every $\omega \in \Omega^\theta$ and $\omega' \in \left(\cup_{\theta' \in \Theta^c}\Omega^{\theta'}\right) \cap B_\epsilon\left(\omega\right)$, there are only three possible scenarios:

1. $\omega \in \Omega^\theta\left(\sigma\right)$ and $K^p\left(a^*, \omega\right) < K^p\left(a^*, \omega'\right)$;

2. $\omega \in \Omega^\theta\left(\sigma\right)$ and $p\left(\cdot|a^*, \omega\right) \equiv p\left(\cdot|a^*, \omega'\right)$;

3. $\omega \notin \Omega^\theta\left(\sigma\right)$ and $K^p\left(a^*, \omega''\right) < K^p\left(a^*, \omega'\right)$ for all $\omega'' \in \Omega^\theta\left(\sigma\right)$.

Notice immediately that in Scenario 2, we have $l_t^{p,\omega'}/l_t^{p,\omega} \equiv 1, \forall t$ on $\tilde{H}$. In Scenario 1, by Lemma 4, $l_t^{p,\omega'}/l_t^{p,\omega} \overset{a.s.}{\to} 0$ on $\tilde{H}$. Similarly in Scenario 3, $l_t^{p,\omega'}/l_t^{p,\omega''} \overset{a.s.}{\to} 0$ on $\tilde{H}$. Therefore, there exists a new full-support $\hat{\pi}_0^\theta$ and a positive measure of paths $\hat{H} \subset \tilde{H}$ where for all t, not only do we have $a_t = a^*, \hat{\pi}_t^\theta \left(\Omega^\theta(\sigma)\right) > \gamma$, but it also holds that $l_t^{p,\omega'}/l_t^{p,\omega} \leq 1$ and $l_t^{p,\omega'}/l_t^{p,\omega''} \leq 1$ for all t and all above combinations of $\omega \in \Omega^\theta$ and $\omega' \in \left(\cup_{\theta' \in \Theta^c}\Omega^{\theta'}\right) \cap B_\epsilon(\omega)$. Now consider any $\theta' \in \Theta^c$ with $\pi_0^{\theta'} \in \Omega^{\theta'} \cup D_\epsilon^p\left(\hat{\pi}_0^\theta\right)$ and the likelihood ratio $l_t^{\theta'}/l_t^\theta$ computed from any history in $\hat{H}$ under prior $\hat{\pi}_0^\theta$,

$$\frac{l_t^{\theta'}}{l_t^\theta} = \frac{\sum_{\omega \in \Omega^{\theta'}} \pi_0^{\theta'}(\omega') \, l_t^{p,\omega'}}{\sum_{\omega \in \Omega^\theta} \hat{\pi}_0^\theta(\omega) \, l_t^{p,\omega}} < \frac{1}{\hat{\pi}_0^\theta(\omega)} < \alpha.$$

Since this is true for each $\theta' \in \Theta^c$, we now know that no switcher will be triggered on $\hat{H}$, thereby completing our proof of p-constrained local robustness.

B.5 Proof of Theorem 4

Suppose that there exist no Berk-Nash equilibrium σ under θ with θ being locally KL-minimizing at σ w.r.t. Ω^p. Besides, let us suppose for the sake of contradiction that θ is p-constrained locally robust within a neighborhood of ϵ.

Take any Berk-Nash equilibrium $\sigma \in \Delta\mathcal{A}$ under θ, then by assumption, there must exist some parameter $\omega' \in \Omega^p$ such that $\min_{\omega \in \Omega^\theta} \|\omega - \omega'\| \leq \epsilon$ and

$$\min_{\omega \in \Omega^\theta} K^p(\sigma, \omega) > K^p(\sigma, \omega'). \tag{16}$$

By continuity, there exists some open neighborhood of σ, denoted as O_σ, in which ω' yields a strictly lower KL divergence than Ω^θ, i.e. $\forall \sigma' \in O_\sigma$, we have

$$\min_{\omega \in \Omega^\theta} K^p(\sigma', \omega) > K^p(\sigma', \omega').$$

We know from Lemma 7 that the set of Berk-Nash equilibria under θ is compact. Therefore, by the Heine-Borel theorem, there must exist finite number of parameters, collected by a set R_ϵ, such that for any Berk-Nash equilibrium σ, we can find some parameter from the set R_ϵ such that the above inequality (16) holds.

Consider a competing model θ^c with an expanded parameter space $\Omega^{\theta^c} = \Omega^\theta \cup R_\epsilon$, and some prior $\pi_0^{\theta^c}$ that allocates a total probability of ϵ evenly to R_ϵ. Formally, let

$$\begin{aligned} \pi_0^{\theta^c}(\omega) &= (1-\epsilon)\,\pi_0^\theta(\omega), \forall \omega \in \Omega^\theta, \\ \pi_0^{\theta^c}(\omega) &= \frac{\epsilon}{|R_\epsilon|}, \forall \omega \in R_\epsilon. \end{aligned}$$

Consider all possible histories in which the switcher eventually adopts θ. Then the switcher's action frequency a.s. converges to a Berk-Nash equilibrium by assumption. Consider the paths where this limit equilibrium is σ. Then it must be that $\limsup_t l_t^{\theta^c}/l_t^{\theta} \leq \alpha$ on those paths. By construction, there exists some $T > 0$ and $\eta > 0$ such that $\forall t > T$, there exists $\omega'' \in R_\epsilon$ such that $K^p(\sigma_t, \omega'') - K^p(\sigma_t, \omega) < -\eta, \forall \omega \in \Omega^\theta$. It then follows that

$$
\begin{aligned}
\lambda_t^{\theta^c} &= \frac{\sum_{\omega' \in \Omega^{\theta^c}} \pi_0^{\theta^c}(\omega') \prod_{\tau=0}^{t-1} q^{\theta^c}(y_\tau | a_\tau, \omega')}{\sum_{\omega \in \Omega^\theta} \pi_0^\theta(\omega) \prod_{\tau=0}^{t-1} q^\theta(y_\tau | a_\tau, \omega)} \\
&> \frac{\frac{\epsilon}{|R_\epsilon|} \prod_{\tau=0}^{t-1} q^{\theta^c}(y_\tau | a_\tau, \omega'')}{\sum_{\omega \in \Omega^\theta} \pi_0^\theta(\omega) \prod_{\tau=0}^{t-1} q^\theta(y_\tau | a_\tau, \omega)} \\
&= \frac{\epsilon}{|R_\epsilon|} \frac{1}{\sum_{\omega \in \Omega^\theta} \pi_0^\theta(\omega) \exp\left(t\left(K^p(\sigma_t, \omega'') - K^p(\sigma_t, \omega)\right)\right)} \\
&> \frac{\epsilon}{|R_\epsilon|} \exp(t\eta)
\end{aligned}
$$

Therefore, for any $\alpha > 0$, almost surely, $l_t^{\theta^c}/l_t^{\theta}$ exceeds α for infinitely many t, contradicting our assumption that $\limsup_t l_t^{\theta^c}/l_t^{\theta} \leq \alpha$ on those paths. Therefore, θ does not persist against θ^c. Since the choice of ϵ is arbitrary, this implies that θ is not p-constrained locally robust.

B.6 Proof of Theorem 5

We only need to show that given any $\epsilon > 0$, almost surely, a dogmatic modeler's action frequency σ_t enters the ϵ-neighborhood of some Berk-Nash equilibrium infinitely often from every full-support prior and policy. Then using a similar argument as in the proof of Theorem 4, it can be shown that θ is not p-constrained locally robust if there is no Berk-Nash equilibrium σ such that θ is locally KL-minimizing at σ.

For convenience, let $\mathcal{A} = \{a^1, a^2\}$. First, consider the paths where σ_t converges to some limit σ. denoted by H^1. Then Lemma 5 tells us that $\pi_t^\theta\left(\Omega^\theta(\sigma)\right)$ converges to 1. Therefore, any action $a \notin \cup_{\pi \in \Delta\Omega^\theta(\sigma)} A_m^\theta(\pi)$ cannot be in the support of σ. Hence, for each action a in the support of σ, there exists some belief $\pi_a \in \Delta\Omega^\theta(\sigma)$ such that $a \in A_m^\theta(\pi_a)$. If $\text{supp}(\sigma)$ is a singleton, then this immediately implies that σ is a Berk-Nash equilibrium. If instead $\text{supp}(\sigma) = \{a^1, a^2\}$, then by the hemi-continuity of A_m^θ, there must exist some $\pi_\sigma \in \Delta\Omega^\theta(\sigma)$ such that $\{a^1, a^2\} = A_m^\theta(\pi_a)$, which again implies that σ is a Berk-Nash equilibrium. Therefore, her action frequency σ_t enters the ϵ-neighborhood of some Berk-Nash equilibrium infinitely often for any $\epsilon > 0$ almost surely on H^1.

Now consider paths where her action frequency oscillates forever, denoted by H^2. Let Ω_∞^θ be the set of all parameters in Ω^θ that are KL minimizers infinitely often, i.e. $\Omega_\infty^\theta = \left\{\omega \in \Omega^\theta : \omega \in \Omega^\theta(\sigma_t) \text{ for infinitely many } t \text{ on } H^2\right\}$. Take any $\omega \in \Omega_\infty^\theta$. Suppose that $A_m^\theta(\delta_\omega) =$

$\{a^1, a^2\}$, then each action frequency σ_ω that satisfies $\omega \in \Omega^\theta(\sigma_\omega)$ is a Berk-Nash equilibrium. By construction, this means σ_t constitutes a Berk-Nash equilibrium infinitely often.

Suppose instead that $\forall \omega \in \Omega^\theta_\infty$, we have that $A^\theta_m(\delta_\omega)$ is singleton. Since σ_t oscillates, Ω^θ_∞ cannot be a singleton. It must be that $A^\theta_m(\delta_\omega) = \{a^1\}$ for some $\omega \in \Omega^\theta_\infty$ and or $A^\theta_m(\delta_{\omega'}) = \{a^2\}$ for some other $\omega' \in \Omega^\theta_\infty$. Given any ω and $\omega' \in \Omega^\theta_\infty$, say they are *related* if there exists some mixed action σ such that $\omega, \omega' \in \Omega^\theta(\sigma)$. I now show that there must exist such a pair of related parameters such that $A^\theta_m(\delta_\omega) = \{a^1\}$ and $A^\theta_m(\delta_{\omega'}) = \{a^2\}$.

First of all, every parameter in Ω^θ_∞ must be related to some other parameter in Ω^θ_∞. Suppose not for the sake of a contradiction. Then there exists some "isolated" parameter $\omega^* \in \Omega^\theta_\infty$ in the following sense: let $C_\omega = \left\{\beta \in [0,1] : \omega \in \Omega^\theta\left(\beta a^1 + (1-\beta)a^2\right)\right\}$, then there exists some positive constant γ such that $B_\gamma(C_{\omega^*}) \cap \left(\cup_{\omega \in \Omega^\theta_\infty \setminus \{\omega^*\}} C_\omega\right) = \emptyset$. However, since ω^* is a KL minimizer infinitely often, it happens infinitely often that $\sigma_t \in C_{\omega^*}$. It implies that some KL minimizer at $\sigma \in B_\gamma(C_{\omega^*}) \setminus C_{\omega^*}$ should also be a KL minimizer at σ_t infinitely often yet not included by Ω^θ_∞, contradicting the definition of Ω^θ_∞. By the same logic, there cannot be two cliques in Ω^θ_∞ such that every parameter in the first clique is unrelated to every parameter in the second clique.

Hence, if every pair of related parameters in Ω^θ_∞ induce the same optimal action, then $A^\theta_m(\delta_\omega) = \{a^1\}$ or $\{a^2\}$ for all $\omega \in \Omega^\theta_\infty$, which we know is not true. Therefore, there exists a related pair $\omega, \omega' \in \Omega^\theta_\infty$ such that $A^\theta_m(\delta_\omega) = \{a^1\}$ and $A^\theta_m(\delta_{\omega'}) = \{a^2\}$. Therefore, each mixed action in $C_\omega \cap C_{\omega'}$ is a Berk-Nash equilibrium. Notice that each C_ω is compact and convex. Since σ_t enters both C_ω and $C_{\omega'}$ infinitely many times, it must be that σ_t enters the ϵ-neighborhood of $C_\omega \cap C_{\omega'}$ infinitely often for any $\epsilon > 0$. The proof is now complete.

B.7 Proof of Theorem 6

B.7.1 Part (i)

For convenience, denote the parameter space of θ as $\Omega^\theta = \left\{\omega^1, ..., \omega^N\right\}$. Suppose that there dose not exist an SCE with a supporting belief π such that $\omega^1 \in \text{supp}(\pi)$.

I now construct a competing model θ^c and later show that θ is not strongly robust against θ^c. Let $\Omega^{\theta^c} = \left\{\omega^*, \omega^2, ..., \omega^N\right\}$ and $q^{\theta^c}(\cdot|\cdot, \omega^n) = q^\theta(\cdot|\cdot, \omega^n)$ for all $n = 2, ..., N$. Also, assume that ω^* corresponds to the objective outcome distribution, i.e. $q^{\theta^c}(\cdot|\cdot, \omega^*) = q^*(\cdot|\cdot)$. Finally, pick any $\epsilon > 0$ and assume that $\pi^\theta_0(\omega^1) = 1 - \pi^{\theta^c}_0(\omega^*) = 1 - \epsilon$ and $\pi^\theta_0(\omega^n)/\pi^\theta_0(\omega^m) = \pi^{\theta^c}_0(\omega^n)/\pi^{\theta^c}_0(\omega^m), \forall m, n = 2, 3, ..., N$.

Suppose θ persists against θ^c at some full-support beliefs and priors. Since θ^c is correctly specified, there must exist a p-absorbing SCE σ under θ. On the paths where m_t converges to θ, the agent eventually only plays actions contained in $\text{supp}(\sigma)$. It follows that $\pi^\theta_t(\omega^1)$

converges to 0 on those paths. Therefore, for any positive $\gamma < 1$, there almost surely exists a stopping time $T > 0$ such that $\pi_t^\theta(\omega^1) < \gamma$ for any $t > T$. Further,

$$\begin{aligned}\frac{l_t^{\theta^c}}{l_t^\theta} &= \frac{(1-\epsilon)\prod_{\tau=0}^{t-1} q^*(y_\tau|a_\tau) + \sum_{n=2}^N \prod_{\tau=0}^{t-1} \pi_0^{\theta^c}(\omega^n)\, q^{\theta^c}(y_\tau|a_\tau,\omega^n)}{(1-\epsilon)\prod_{\tau=0}^{t-1} q^\theta(y_\tau|a_\tau,\omega^1) + \sum_{n=2}^N \prod_{\tau=0}^{t-1} \pi_0^\theta(\omega^n)\, q^\theta(y_\tau|a_\tau,\omega^n)} \\ &> \frac{(1-\epsilon)/\epsilon}{\pi_t^\theta(\omega^1)/\left(1-\pi_t^\theta(\omega^1)\right)+1} \\ &> \frac{(1-\epsilon)/\epsilon}{\gamma/(1-\gamma)+1}.\end{aligned}$$

Let ϵ be small enough so that $\frac{(1-\epsilon)/\epsilon}{\gamma/(1-\gamma)+1} > \alpha$. This implies that a switch is eventually triggered, contradicting the assumption that m_t converges to θ with positive probability. Therefore, θ is not strongly robust, as desired.

B.7.2 Part (ii)

I first show that given any discount factor $\delta \leq 1$, the agent only chooses actions from the set $A_m^\theta(\Delta\Omega^\theta)$. Given the current belief π_t^θ, action $a' \in \mathcal{A}$ maximizes the period-t flow payoff under θ if

$$\begin{aligned} a' &\in \arg\max_a \sum_{\omega\in\Omega^\theta} \pi_t^\theta(\omega) \int_{\mathcal{Y}} u(a,y)\, q^\theta(y|a,\omega)\, v(dy) \\ &= \arg\max_a \int_{\mathcal{Y}} u(a,y)\, q^*(y|a)\, v(dy), \end{aligned}$$

where the equality follows from the assumption that $q^\theta(\cdot|a,\omega) \equiv q^*(\cdot|a), \forall\omega$ for all $a \in A_m^\theta(\Delta\Omega^\theta)$. Notice that the set of myopically optimal actions does not vary with π_t^θ. Therefore, there is no experimentation value for any action outside $A_m^\theta(\Delta\Omega^\theta)$.

Suppose that $q^\theta(\cdot|a,\omega) \equiv q^*(\cdot|a)$ for all $a \in A^\theta(\Delta\Omega^\theta)$ and all $\omega \in \Omega^\theta$. Since $a_\tau \in A^\theta(\Delta\Omega^\theta)$ for all τ, the likelihood ratio of θ^c to θ

$$\frac{l_t^{\theta^c}}{l_t^\theta} = \frac{\sum_{\omega\in\Omega^{\theta^c}} \pi_0^{\theta^c}(\omega) \prod_{\tau=0}^{t-1} q^{\theta^c}(y_\tau|a_\tau,\omega)}{\prod_{\tau=0}^{t-1} q^*(y_\tau|a_\tau)}$$

is a positive supermartingale, with $l_0^{\theta^c}/l_0^\theta = 1$. By the maximal inequality, for any $\gamma > 0$,

$$\mathbb{P}^*\left(l_t^{\theta^c}/l_t^\theta \geq 1+\gamma \text{ for some } t \geq 0\right) < 1.$$

Therefore, for any fixed $\alpha > 1$, the probability that $l_t^{\theta^c}/l_t^\theta$ never exceeds α is strictly positive, thereby implying that θ is strongly robust.

References

Akaike, H. (1974): "A New Look at the Statistical Model Identification," *IEEE Transactions on Automatic Control*, 19, 716–723.

Ba, C. and A. Gindin (2021): "A Multi-Agent Model of Misspecified Learning with Overconfidence," *Available at SSRN*.

Berk, R. H. (1966): "Limiting Behavior of Posterior Distributions When the Model is Incorrect," *The Annals of Mathematical Statistics*, 51–58.

Blackwell, D. (1965): "Discounted Dynamic Programming," *The Annals of Mathematical Statistics*, 36, 226–235.

Bohren, J. A. (2016): "Informational Herding with Model Misspecification," *Journal of Economic Theory*, 163, 222–247.

Bohren, J. A. and D. N. Hauser (2021): "Learning with Heterogeneous Misspecified Models: Characterization and Robustness," *Econometrica, forthcoming.*

Braungart, R. G. and M. M. Braungart (1986): "Life-course And Generational Politics," *Annual Review of Sociology*, 12, 205–231.

Chernoff, H. (1954): "On the Distribution of the Likelihood Ratio," *The Annals of Mathematical Statistics*, 573–578.

Cho, I.-K. and K. Kasa (2015): "Learning and Model Validation," *The Review of Economic Studies*, 82, 45–82.

Di Stefano, G., F. Gino, G. P. Pisano, and B. Staats (2015): *Learning by Thinking: Overcoming the Bias for Action through Reflection*, Harvard Business School Cambridge, MA, USA.

Easley, D. and N. M. Kiefer (1988): "Controlling a Stochastic Process with Unknown Parameters," *Econometrica: Journal of the Econometric Society*, 1045–1064.

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

Esponda, I., D. Pouzo, and Y. Yamamoto (2019): "Asymptotic Behavior of Bayesian Learners with Misspecified Models," *arXiv preprint arXiv:1904.08551*.

Eyster, E. and M. Rabin (2010): "Naive Herding in Rich-information Settings," *American Economic Journal: Microeconomics*, 2, 221–43.

Fisher, W. R. (1985): "The Narrative Paradigm: An Elaboration," *Communications Monographs*, 52, 347–367.

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

——— (2021): "Welfare Comparisons for Biased Learning," *Cowles Foundation Discussion Papers. 2605.*

Fudenberg, D. and G. Lanzani (2020): "Which Misperceptions Persist?" *Available at SSRN.*

Fudenberg, D., G. Lanzani, and P. Strack (2021): "Limit Points of Endogenous Misspecified Learning," *Econometrica*, 89, 1065–1098.

Fudenberg, D. and D. K. Levine (1993): "Self-Confirming Equilibrium," *Econometrica*, 523–545.

Fudenberg, D., G. Romanyuk, and P. Strack (2017): "Active Learning with a Misspecified Prior," *Theoretical Economics*, 12, 1155–1189.

Gagnon-Bartsch, T. (2017): "Taste Projection in Models of Social Learning," Tech. rep., Mimeo.

Gagnon-Bartsch, T., M. Rabin, and J. Schwartzstein (2020): *Channeled Attention and Stable Errors*, Harvard Business School.

Galperti, S. (2019): "Persuasion: The Art of Changing Worldviews," *American Economic Review*, 109, 996–1031.

He, K. (2020): "Mislearning from Censored Data: The Gambler's Fallacy in Optimal-Stopping Problems," *arXiv preprint arXiv:1803.08170.*

He, K. and J. Libgober (2020): "Evolutionarily Stable (Mis) specifications: Theory and Applications," *arXiv preprint arXiv:2012.15007.*

Heidhues, P., B. Kőszegi, and P. Strack (2018): "Unrealistic Expectations and Misguided Learning," *Econometrica*, 86, 1159–1214.

——— (2019): "Overconfidence and Prejudice," *arXiv preprint arXiv:1909.08497.*

Karni, E. and M.-L. Vierø (2013): ""Reverse Bayesianism": A Choice-based Theory Of Growing Awareness," *American Economic Review*, 103, 2790–2810.

Kass, R. E. and A. E. Raftery (1995): "Bayes Factors," *Journal of the american statistical association*, 90, 773–795.

Konishi, S. and G. Kitagawa (2008): *Information Criteria and Statistical Modeling*, Springer Science & Business Media.

Kuhn, T. S. (1962): *The Structure of Scientific Revolutions*, University of Chicago Press.

Mailath, G. J. and L. Samuelson (2020): "Learning under Diverse World Views: Model-Based Inference," *American Economic Review*, 110, 1464–1501.

Maitra, A. (1968): "Discounted Dynamic Programming on Compact Metric Spaces," *Sankhyā: The Indian Journal of Statistics, Series A*, 211–216.

Massey, C. and G. Wu (2005): "Detecting Regime Shifts: The Causes of Under-and Overreaction," *Management Science*, 51, 932–947.

Mullainathan, S. (2002): "Thinking Through Categories," Tech. rep., Working Paper, Harvard University.

Nyarko, Y. (1991): "Learning in Mis-specified Models and the Possibility of Cycles," *Journal of Economic Theory*, 55, 416–427.

Olea, J. L. M., P. Ortoleva, M. M. Pai, and A. Prat (2019): "Competing Models," *arXiv preprint arXiv:1907.03809*.

Ortoleva, P. (2012): "Modeling the Change of Paradigm: Non-bayesian Reactions to Unexpected News," *American Economic Review*, 102, 2410–36.

Ortoleva, P. and E. Snowberg (2015): "Overconfidence in Political Behavior," *American Economic Review*, 105, 504–35.

Rabin, M. and D. Vayanos (2010): "The Gambler's and Hot-hand Fallacies: Theory and Applications," *Review of Economic Studies*, 77, 730–778.

Savage, L. J. (1972): *The Foundations of Statistics*, Courier Corporation.

Schwartzstein, J. and A. Sunderam (2019): "Using Models to Persuade," Tech. rep., National Bureau of Economic Research.

Schwarz, G. et al. (1978): "Estimating the Dimension of a Model," *The Annals of Statistics*, 6, 461–464.

Spiegler, R. (2016): "Bayesian Networks and Boundedly Rational Expectations," *The Quarterly Journal of Economics*, 131, 1243–1290.

——— (2019): "Behavioral Implications of Causal Misperceptions," *Annual Review of Economics*, 12.

——— (2020): "Can Agents with Causal Misperceptions be Systematically Fooled?" *Journal of the European Economic Association*, 18, 583–617.

Stone, M. (1977): "An Asymptotic Equivalence of Choice of Model by Cross-validation and Akaike's Criterion," *Journal of the Royal Statistical Society: Series B (Methodological)*, 39, 44–47.

Ville, J. (1939): "Ètude Critique de la Notion de Collectif," *Bulletin of the American Mathematical Society*, 45, 824.

Wegener, D. T. and R. E. Petty (1997): "The Flexible Correction Model: The Role of Naive Theories of Bias in Bias Correction," in *Advances in experimental social psychology*, Elsevier, vol. 29, 141–208.

C Online Appendix

I provide an example below to substantiate the observation in Footnote 5.

Example 6. Let x^1 and x^2 be two i.i.d. normally distributed variables, both with mean 0 and variance 1. Suppose x^3 and x^4 are also i.i.d. normally distributed but with mean 1 and variance 1. Suppose the agent can play one of two actions in each period, $\mathcal{A} = \{1, 2\}$ and uses subjective models to learn about the mean of each element in y. Her flow payoff is given by $a \cdot (x^4 - x^3)$. Hence, she would like to play $a = 2$ if $\overline{x}_4 > \overline{x}_3$ and play $a = 1$ if $\overline{x}_3 > \overline{x}_4$. However, x_1 and x_3 are only observable when $a = 1$, while x_2 and x_4 are only observable when $a = 2$. She entertains an initial model θ and two competing models, $\{\theta^1, \theta^2\}$, each of which is equipped with a binary parameter space. The predictions of each model are summarized by the following table.

θ	ω^1	ω^2
$(\overline{x}^1, \overline{x}^2, \overline{x}^3, \overline{x}^4)$	$(1, 1, 1, 0)$	$(1, 1, 0, 1)$
θ^1	$\omega^{1\prime}$	$\omega^{2\prime}$
$(\overline{x}_1, \overline{x}_2, \overline{x}_3, \overline{x}_4)$	$(1, 0, 1, 0)$	$(1, 0, 0, 1)$
θ^2	$\omega^{1\prime\prime}$	$\omega^{2\prime\prime}$
$(\overline{x}_1, \overline{x}_2, \overline{x}_3, \overline{x}_4)$	$(0, 1, 1, 0)$	$(0, 1, 0, 1)$

Notice that there are two uniformly strict and thus p-absorbing Berk-Nash equilibria under θ: (1) $a = 1$ is played w.p. 1, supported by the belief that assigns probability 1 to ω^1; (2) $a = 2$ is played w.p. 1 , supported by the belief that assigns probability 1 to ω^2. First observe that θ persists against θ^1 at a prior π_0^θ that assigns sufficiently high belief to ω^1. This follows from the fact that the likelihood ratio between θ and θ^1 is always 1 when $a = 1$ is played, and that the equilibrium is p-absorbing. Analogously, θ persists against θ^2 at a prior π_0^θ that assigns sufficiently high belief to ω^2. However, notice that θ does not persist against $\{\theta^1, \theta^2\}$ at any priors and policies, because regardless of the actions taken by the agent, at least one of θ^1 and θ^2 would fit the data strictly better than θ, prompting the agent to adopt θ^1 and θ^2 infinitely often.

Example 7 below shows that a model can fail to be constrained locally robust even if it induces a p-absorbing BN-E satisfying the local KL-minimization property but not the local identification property.

Example 7. Let $\mathcal{Y} = \mathbb{R}^2$ and $\mathcal{A} = \{a\}$. Denote the two-dimensional outcome y as (y_1, y_2) and assume that y_1 and y_2 are independently and normally distributed with zero mean, i.e.

$y \sim N\left((0,0)', I\right)$. Consider a family of normal densities $p(y|a,\omega)$ where $\omega = (\omega_1, \omega_2) \in \Omega^\theta \subset \mathbb{R}^2$ corresponds to a joint standard normal distribution with mean ω and covariance I. Specifically, let $\Omega = \{(\omega_1, \omega_2) \in \mathbb{R}^2 : \omega_1^2 + \omega_2^2 = 1\}$ and $\Omega^\theta = \{(1,0)\}$. Notably, $K^p(a,\omega)$ is a constant over Ω^p. Hence, θ is locally KL-minimizing but not locally identified at any σ w.r.t. Ω^p. Fix any $\epsilon > 0$, we can find a competing model with $\Omega^{\theta^c} = \{(\cos\phi, \sin\phi)\}$, where ϕ is close enough to 0 so that the distance between Ω^{θ^c} and Ω^θ is strictly smaller than ϵ. However, the agent will adopt θ^c infinitely often as the log-likelihood ratio is an unbounded random walk. Intuitively, since θ and θ^c predict equally well in the long term, it must happen infinitely many times that some streak of outcomes increases the likelihood ratio to above α and trigger a switch to θ^c if the agent adopted θ.