# Robust Model Misspecification and Paradigm Shifts

Cuimin Ba*

First version: Apr 14, 2021

This version: Mar 5, 2022

– Click here to see the most recent version –

**Abstract**

This paper studies the forms of model misspecification that are more likely to persist when an agent compares her subjective model with competing models. The agent learns about an action-dependent outcome distribution and makes decisions repeatedly. 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 competing model and is *locally robust* if it can persist against every nearby competing model under nearby priors. The main result provides a simple characterization of globally robust and locally robust models based on the set of Berk-Nash equilibria they induce. I then use these results to provide the first learning foundations for the persistence of systemic biases in two canonical applications.

**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, George Mailath, and Kevin He for their guidance and support at every stage of this paper. I thank Alvaro Sandroni, 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 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 act suboptimally.[1]

The assumption that individuals forever hold on to a single misspecified model contradicts a plethora of evidence. Economists and data scientists, for example, often switch their models of the world when an alternative seems to better fit the data: they may add new explanatory variables to a regression equation, or adopt the Natural Rate Hypothesis in place of the Phillips Curve. The history of science offers numerous examples of paradigm shifts in scientific advances (Kuhn, 1962). More generally, people may have multiple subjective models and switch from one to another. Individuals 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 abandon their current model? Which forms of misspecification are more likely to persist? In other words, which subjective models are *robust*? This paper addresses these questions using a framework of misspecified Bayesian learning that allows agents to revise their models.

I consider an agent facing an infinite-horizon decision problem. In each period, the agent chooses an action and then observes an outcome, the true distribution of which is unknown and contingent on the action. The agent then obtains a flow payoff jointly determined by the action and the realized outcome. A subjective model is a theory of how actions affect the random outcomes. In contrast to a *dogmatic modeler* who relies on a single model, I consider a *switcher* who switches between multiple models. In the baseline framework, she starts with an *initial* model, while simultaneously considering a single *competing* model as a potential replacement. The models are parametric: a model, together with a specific

[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).

parameter value, corresponds to a profile of action-dependent outcome distributions. In each period, she either keeps the model she used last period or switches to the alternative. The agent uses the current model to complete two tasks: she updates 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 conceptual 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 a subjective model to interpret the world and guide her actions, until the data reveal to her the superiority of a competing model. Within a paradigm, she behaves exactly as prescribed by the theory of subjective probability (Savage, 1972), updating beliefs using Bayes rule and behaving in a dynamically consistent manner. In contrast, a paradigm shift reflects fundamental changes to the agent's subjective assumptions about the world, which cannot be captured by a Bayesian approach that assigns a prior to a set of models. I assume that the agent switches to a different model only if it is compelling enough according to a *Bayes factor criterion* (Kass and Raftery, 1995). The Bayes factor of 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 if the likelihood ratio is above an exogenous switching threshold that is larger than 1 and makes no switch otherwise.

A model is said to *persist* against a given competing model if on a positive measure of paths, when starting at this model, the agent eventually adopts it forever, potentially after many rounds of switching. The main analysis is concerned with the properties of the model that lead to persistence. Several challenges arise when incorporating a model switching process into a misspecified learning problem. Given that the agent has access to multiple models, we need to keep track of multiple belief processes, all of which 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 literature, such belief processes and the implied sequence of actions may oscillate forever if the models are misspecified. Consequently, the Bayes factor may fail to have good convergence properties. Moreover, the agent's best response function changes with her model choice, intertwining the multiple belief processes and making it difficult to assess the persistence of models.

I develop three different notions of robustness to measure the persistence of models. A model is *globally robust* if it can persist against every possible competing model. This notion of robustness requires persistence no matter what competing model is entertained. When the agent has a limited understanding of the world or when she is reluctant to consider significant changes, however, it makes sense to restrict attention to competing models that are close to the initial model. I say that a model is *unconstrained locally robust* if this model can persist against every possible "neighbor" model with similar priors. Here, models are nearby each other if they predict similar outcome distributions. This notion is unconstrained because it places no restriction on the similarity of the parametric structures. In contrast, a model is *constrained locally robust* if it can persist against every nearby model from the same parametric family under similar priors.[2] These robustness notions allow us to abstract from specifying a particular competing model and lead to clean characterizations of misspecification persistence based on the primitives of the misspecified model.

The main results of this paper characterize each notion of robustness. It is unclear whether a globally robust model even exists given its strong requirements. I first provide a characterization of global robustness and show that even misspecified models can be globally robust. Since a globally robust model can, by definition, persist against a correctly specified model, it should have perfect goodness-of-fit at least in the limit. In a *self-confirming equilibrium*, the agent holds a belief that exactly matches the objective outcome distribution and plays a strategy optimal against this belief. I say that a self-confirming equilibrium under a model is *p-absorbing* if the action of a dogmatic modeler who only uses this model converges to the support of the equilibrium with positive probability. Theorem 1 establishes that the existence of such an equilibrium is necessary and sufficient for global robustness. This equivalence reduces the complicated problem of a switcher to the problem of a dogmatic modeler. Building on existing results from the literature, Corollary 2 provides a sufficient condition for the p-absorbing condition which is straightforward to verify from the primitives.

The intuition behind Theorem 1 is as follows: when the agent starts with a prior that is close to a self-confirming 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 with positive probability, she never deviates from the equilibrium. On those paths, a switch is never triggered, as the agent already has entrenched beliefs and plays a self-confirming equilibrium. But what happens when the agent has not yet reached such a belief? Theorem 2 shows that entrenched priors are in fact necessary. With a more diffuse prior, a competing model can be constructed such that the agent is guaranteed to abandon her initial model.

I next turn to the characterization of local robustness. Theorem 3 shows that unconstrained local robustness is equivalent to global robustness. 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

[2] For example, for a monopolist who is trying to estimate a market demand function that he believes to be linear, unconstrained local robustness would require persistence against every nearby demand function, potentially complicated and non-linear, while the constrained notion only requires persistence against every nearby linear function.

generating process. Constrained local robustness, on the other hand, is much weaker than global robustness. Theorem 4 establishes that a model is constrained locally robust if there exists a pure p-absorbing Berk-Nash equilibrium that satisfies an additional property called *local dominance*. This property plays a similar role as the self-confirming condition in the case of global robustness and ensures that when the Berk-Nash equilibrium is being played, no local perturbation of the model can fit the data strictly better. Theorem 5 focuses on the necessary conditions for constrained local robustness and show that in some special environments, constrained local robustness requires the existence of a Berk-Nash equilibrium at which the model can yield a weakly lower Kullback-Leibler divergence than any neighbor models from the same family.

In an extension, I investigate how the characterizations might change when the agent entertains multiple competing models at the same time. I find that, due to overfitting, even the true model may fail to be persistent if the number of competing models is unrestricted. Nevertheless, if the agent is constrained to evaluate at most $K$ competing models and switching is sticky enough in that the switching threshold is high relative to $K$, the previous characterizations still apply (Theorem 6).

I use these results to study the persistence of misspecification in two applications. The first application studies a financial investment problem and shows that extreme pessimism about the stock market is more likely to persist than extreme optimism. Extremely pessimistic investors play a self-confirming equilibrium by staying away from the market and resorting to safe outside options, while extremely optimistic investors get involved in trading and gradually learn about the true market condition.

The second application studies a variation of Heidhues et al. (2018) in which an individual with an incorrect self-perception makes effort choices and may attribute the discrepancy in observed and expected outputs to an outside fundamental, such as his teammate's ability. The main result establishes that overconfidence is more likely to be robust than underconfidence, consistent with the mounting empirical evidence that individuals tend to exhibit overconfidence on average. I demonstrate that overconfidence causes the individual's belief about the fundamental to be positively reinforcing while underconfidence makes it negatively reinforcing, which have qualitatively different implications on the individual's equilibrium behavior and thus the robustness of the biases. This result provides a novel mechanism that breaks the symmetry between overconfidence and underconfidence, and shows that the prevalence of overconfidence may be a result of the learning process itself, rather than external factors such as ego utility from positive beliefs.

The rest of the paper is organized as follows. The next subsection discusses related literature. Section 2 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 that allows for multiple competing models. Section 6 presents the applications and Section 7 concludes. Appendix A contains useful auxiliary results, and Appendix B includes all proofs of the main results.

## Related Literature

This paper builds on the literature of learning with misspecified models.[3] Esponda and Pouzo (2016) first proposed the concept of Berk-Nash equilibrium. 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 divergence minimizers, the dogmatic modeler's action converges to the equilibrium with arbitrarily high probability. Frick et al. (2021a) introduce a prediction accuracy ordering over parameters that is different from the ordering based on Kullback-Leibler divergence and use martingale convergence arguments to prove that the belief process of the dogmatic modeler converges to a point belief at the top ranked parameter. My paper contributes to the literature by allowing for model switching and proposing various robustness notions for misspecified models to persist. Since the persistence of a model partly hinges on the asymptotic behavior of a dogmatic modeler who holds this model throughout, some technical challenges in this paper relate to those in the literature. On the other hand, the Bayes factor process that governs model switching interacts and correlates with all belief processes (even when no switching has been made), creating new challenges unique to this framework.

This paper is most related to a recent set of papers that explore 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-

[3]Examples include 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 (2021a), all of which study asymptotic learning outcomes of dogmatic modelers in relatively general environments.

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 become more prevalent. 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 (2021b) 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.

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. This consideration is not present in my framework as all subjective distributions and the objective distribution have full support. Karni and Vierø (2013) provide a choice-based decision theory to model a self-correcting agent who can expand his universe of subjective states. 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 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 considering 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 an Euclidean space or a compact subset of an Euclidean space. 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 $\nu$, and the Radon-Nikodym derivative $q^* (\cdot|a)$ is positive; (ii) For all $a \in \mathcal{A}$, $u (a, \cdot) \in L^1 (\mathcal{Y}, \mathbb{R}, Q^* (\cdot|a))$.*[4]

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 $\theta$, consists of two components: (1) a parameter set $\Omega^\theta$ 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 $\Omega^\theta$. I restrict attention to subjective models that satisfy Assumption 2.

**Assumption 2.** *For all $a \in \mathcal{A}$: (i) $\Omega^\theta$ 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 $\nu$, 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 $\left|\ln \frac{q^*(\cdot|a)}{q^\theta(\cdot|a,\omega)}\right| \le g_a (\cdot)$ a.s.-$Q^* (\cdot|a)$.*

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

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 that the difference 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 $\Theta$ be the set of all models $\theta$ that satisfy Assumption 2. Since each element of $\Theta$ is a finite vector of conditional distributions, we have $\Theta \subseteq \cup_{z=1}^{\infty} (\Delta\mathcal{Y})^{|\mathcal{A}|z}$, where $z$ represents the size of the parameter set. A model $\theta$ 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 of $\theta$ includes the true DGP, and *misspecified* otherwise.

## 2.3 The Switcher's Problem

The agent has access to a finite set of subjective models, $\Theta^\dagger \subseteq \Theta$. 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 two subjective models, $\Theta^\dagger = \{\theta^0, \theta^c\}$ (the setting is extended to allow for more than two subjective models in Section 5). 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* $\theta^0$, while evaluating a *competing model* $\theta^c$. 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*. At the beginning of each period $t \geq 1$, the agent calculates a vector of likelihood ratios, or *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, $\lambda_t^\theta$ is the ratio of the likelihood of model $\theta$ to the likelihood of the last period's model choice $m_{t-1}$, evaluated at the initial priors. Let $\theta^* := \arg\max_{\theta\in\Theta^\dagger} \lambda_t^\theta$. If $\lambda_t^{\theta^*} > \alpha$, then a switch to $\theta^*$ 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.[5]

Essentially, the agent is conducting a thought experiment: had I adopted the alternative model, would it better explain the observations? As $\alpha$ becomes larger, switching requires stronger evidence. Thus, $\alpha$ 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 motivation for the assumption that $\alpha > 1$ in Section 2.5.

**Learning.** After pinning down the model, the agent updates her belief over the parameters of $m_t$ 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. With the belief processes, we can equivalently write the likelihood ratios in (1) as

$$\lambda_t^\theta = \lambda_{t-1}^\theta \times \frac{\sum_{\omega \in \Omega^\theta} \pi_{t-1}^\theta(\omega)\, q^\theta(y_{t-1}|a_{t-1}, \omega)}{\sum_{\omega' \in \Omega^{m_{t-1}}} \pi_{t-1}^{m_{t-1}}(\omega)\, q^{m_{t-1}}(y_{t-1}|a_{t-1}, \omega')}. \tag{4}$$

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

$$U^\theta\left(\pi^\theta\right) = \max_{a \in \mathcal{A}} \sum_{\omega \in \Omega^\theta} \pi^\theta(\omega) \int_{y \in \mathcal{Y}} \left[u(a, y) + \delta U^\theta\left(B^\theta\left(a, y, \pi^\theta\right)\right)\right] q^\theta(y|a, \omega)\, v(dy). \tag{5}$$

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, i.e. the solution to (5) 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 $\mathscr{Y} := (\mathcal{Y}^\infty)^{\mathcal{A}}$, each element of which consists of infinite sequences of outcome realizations $(y_{a,0}, y_{a,1}, \ldots)$ for all actions $a \in \mathcal{A}$, where $y_{a,\tau}$ denotes the outcome when the agent takes

[5]This symmetry in the switching threshold is made to simplify notation. All results remain unchanged if the agent uses different thresholds for different switches.

$a \in \mathcal{A}$ in period $\tau$. Denote by $\mathbb{P}$ the probability measure over $\mathscr{Y}$ 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 $\alpha$, the set of models $\Theta^\dagger$, the initial model $\theta^0$, the priors and policies $\left(\pi_0^\theta, a_0^\theta\right)_{\theta \in \Theta^\dagger}$, $\mathbb{P}$ induces a probability measure over $H$ when the agent is a switcher, denoted by $\mathbb{P}_S$. Meanwhile, $\mathbb{P}$ and $\left(\pi_0^\theta, a_0^\theta\right)$ induce a different probability measure over $H$ when the agent is the dogmatic modeler who believes in $\theta$, denoted by $\mathbb{P}_B$.

## 2.4 Persistence of Models

We are interested in situations where the agent eventually settles down with the initial model. For convenience, given any set of models $\Theta' \subseteq \Theta$, 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 $\theta$ is said to *persist in* $\Theta^\dagger := \{\theta, \theta^c\}$, or *persist against* $\theta^c$ 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 $\theta$ as the initial model, the model choice $m_t$ converges to $\theta$ with positive probability.

Persistence will play a crucial role in our definition of model robustness. If $\theta$ does not persist against $\theta^c$, the competing model $\theta^c$ must be adopted by the switcher infinitely many times. As a consequence, 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 $\theta$ and the true DGP. By contrast, if $\theta$ persists, then with positive probability, the agent switcher's behavior resembles the behavior of a dogmatic modeler in the long term. The characterization of persistence not only depends on the structure of the competing model, but may also be affected by other primitives including the prior and policy adopted for each model, the agent's discount factor $\delta$, and her switching threshold $\alpha$. By requiring a model to persist against a number of competing models and varying the restrictions imposed on the competing models and the primitives, we obtain different notions of model 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 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. Therefore, model switching features stickiness, potentially due to the physical or mental cost of discovering and shifting paradigms. 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. Since she chooses actions presuming that no switch will happen in the future, 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 estimate his model is hard enough, so he does not spend 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.

**Switching rule.** The Bayes factor rule enjoys a few advantages among the 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. Lastly, the Bayes factor rule automatically includes a penalty for including too much structure into the model and thus helps prevent overfitting. This is manifested in the comparison with a switching rule that evaluates the likelihood ratio between models using the maximum likelihood estimate of the parameters instead of calculating a weighted sum of likelihoods across all parameters using the prior as weights. This gives an advantage to models that fit better in the short term but are more complicated in the sense that they have larger parameter sets and include more profiles of outcome distributions. The last paragraph in Section 5 comes back to this discussion and further illustrates how the Bayes factor rule guards against overfitting.

## 3 Berk-Nash Equilibria

Our main question concerns when a model $\theta$ persists against a competing model. By definition, if $\theta$ persists, then there exists a first period $T$ such that, with positive probability, the switcher adopts $\theta$ and never switches to $\theta^c$ thereafter. From there, her actions will be identical to a dogmatic modeler who shares the same policy and belief over $\Omega^\theta$ at the onset of period $T$. Hence, whether the switcher can hold on to $\theta$ is closely related to how a dogmatic modeler behaves.

Characterizing the asymptotic behavior of a dogmatic modeler is an important question in 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

$$D_{KL}(q \parallel q') := \int_{\mathcal{Y}} q \ln(q/q') \, \nu(dy). \tag{6}$$

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, for any strategy $\sigma$, let

$$\Omega^\theta(\sigma) := \arg\min_{\omega' \in \Omega^\theta} \sum_{\mathcal{A}} \sigma(a) \, D_{KL}\left(q^*(\cdot|a) \parallel q^\theta(\cdot|a, \omega')\right). \tag{7}$$

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

**Definition 2.** Strategy $\sigma \in \Delta\mathcal{A}$ is a *Berk-Nash equilibrium* (BN-E) under $\theta$ if there exists a belief $\pi \in \Delta\Omega^\theta(\sigma)$ with $\sigma \in \Delta A_m^\theta(\pi)$. A BN-E $\sigma$ 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 $\text{supp}(\sigma) \subseteq A_m^\theta(\pi)$ and $q^*(\cdot|a) \equiv q^\theta(\cdot|a, \omega)$ for all $\omega \in \text{supp}(\pi)$ and all $a \in \text{supp}(\sigma)$.

A Berk-Nash equilibrium requires myopic optimality against a belief $\pi$ that takes support on the KL minimizers; both $\sigma$ and $\pi$ could be non-degenerate. Quasi-strictness further requires that the support of $\sigma$ includes all myopically optimal actions at $\pi$. 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 $\sigma$. Analogously, a BN-E is *uniformly strict* if it is pure and uniformly quasi-strict. Finally, a self-confirming equilibrium requires that each parameter in the support of $\pi$ predicts the same outcome distribution as the true DGP at each action $a$ that is played with positive probability. A *quasi-strict SCE* and a *uniformly quasi-strict SCE* can be defined similarly. While every subjective model admits at least one Berk-Nash equilibrium (Esponda and Pouzo, 2016), the existence of a self-confirming equilibrium is not guaranteed.

# 4 Robustness

## 4.1 Global Robustness

Our notions of robustness build on the concept of persistence. I first define *global robustness*, which requires a model to persist against *every* possible competing model.

**Definition 3** (Global robustness)**.** A model $\theta$ is *globally robust* if there exists a full-support $\pi_0^\theta$ and an optimal $a^\theta$ under which $\theta$ persists against every $\theta^c \in \Theta$ at every full-support $\pi_0^{\theta^c}$ and optimal $a^{\theta^c}$.

Global robustness ensures that $\theta$ can persist no matter what alternative model it is compared against (thus global), as long as the agent starts with a proper choice of prior and policy for $\theta$. Conversely, if $\theta$ is not globally robust, then no matter what prior and policy are employed, one can find a competing model associated with some prior and some policy so that it is almost surely (a.s.) adopted infinitely often.

Let's start with deriving a necessary condition for global robustness. Intuitively, the better the competing model $\theta^c$ is at predicting the outcome distributions, the harder it is for an initial model $\theta$ to persist. So a natural candidate for consideration is a correctly specified competing model. As Lemma 4 in Appendix A shows, if $\theta^c$ is correctly specified, the likelihood ratio of $\theta$ to $\theta^c$, or $l_t^\theta / l_t^{\theta^c}$, is a martingale that a.s. converges. Hence, on paths where $\theta$ is eventually forever adopted, the inverse likelihood ratio $l_t^{\theta^c}/l_t^\theta$—which eventually equals the Bayes factor $\lambda_t^{\theta^c}$—a.s. converges to some value below $\alpha$. Note that since $\theta^c$ is correctly specified, the associated posterior $\pi_t^{\theta^c}$ must converge (regardless of the action history); but this means that the posterior $\pi_t^\theta$ must also stabilize, as Equation (4) tells us that any oscillation of beliefs would imply oscillation of the likelihood ratio as well, contradicting the previous observation. We thus have the following lemma.

**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 $a^\theta, a^{\theta^c}$. Then on paths where $m_t$ converges to $\theta$, almost surely, $l_t^{\theta^c}/l_t^\theta$ converges to a random variable $\iota \leq \alpha$, and $\pi_t^{\theta^c}$ converges to a random variable $\pi_\infty^{\theta^c} \in \Delta\Omega^{\theta^c}$, $\pi_t^\theta$ converges to a random variable $\pi_\infty^\theta \in \Delta\Omega^\theta$.*

An implication of Lemma 1 is that $\theta$ must be able to perfectly predict the distribution of outcomes in the limit. 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 $\theta^c$ persistently outperforms $\theta$ 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, $\pi_t^\theta$ must also assign probability approaching 1 to the true outcome distribution in the limit on a positive measure of paths.

Since data is endogenously generated, this further requires that the agent switcher, potentially after a lot of switches and belief updating, 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, with positive probability, end up playing *only* the equilibrium actions. Since on paths where $\theta$ is adopted forever, a switcher eventually behaves no differently than a dogmatic modeler, this condition can be formalized as a property of the dogmatic modeler as follows.

**Definition 4.** A BN-E $\sigma \in \Delta\mathcal{A}$ under $\theta$ is *absorbing with positive probability*, or *p-absorbing* if under some full-support $\pi_0^\theta$ and optimal $a^\theta$, there exists $T \geq 0$ such that, with positive probability, a dogmatic modeler with $\theta$ only plays actions in $\text{supp}(\sigma)$ after period $T$.

That a BN-E with support $A \subseteq \mathcal{A}$ 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.[6] 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.[7] Example 1 below shows a self-confirming equilibrium that is not p-absorbing.

[6]For example, this is weaker than the stability notion proposed by Fudenberg et al. (2021). They define that a pure BN-E $a^*$ under $\theta$ is stable if for every $\kappa \in (0, 1)$, there exists a belief $\pi \in \Delta\Omega^\theta$ such that for any prior $\pi_0^\theta$ sufficiently close to $\pi$, the dogmatic modeler's action sequence $a_t$ converges to $a^*$ with probability larger than $\kappa$. They do not define a stability notion for mixed BN-E.

[7]Example 5 in Appendix C shows that the agent's action may never converge when she plays a p-absorbing SCE.

**Example 1** (A self-confirming equilibrium that fails to be p-absorbing)**.** Suppose there are two actions $\mathcal{A} = \{1, 3\}$ and three parameters $\Omega^\theta = \{1, 2, 3\}$ inside the parameter space of model $\theta$. The agent's payoff is the absolute value of the outcome, $|y_t|$, with the true DGP of $y_t$ being a normal distribution $N(1, 1)$. Consider a misspecified model $\theta$ that predicts $y_t \sim N(\omega - a_t, 1)$. Note that $\theta$ admits a single self-confirming equilibrium in which the agent plays $a^* = 1$ with probability 1, supported by a belief that assigns probability 1 to $\omega^* = 2$. However, this SCE is not p-absorbing. To see that, notice that the agent is indifferent between the two actions when the parameter takes the value of 2. When the agent keeps playing $a = 1$, the parameters 1 and 3 fit the data equally well on average, so their log-posterior ratio is a random walk which a.s. crosses 1 infinitely often. However, the high action $a = 3$ is strictly optimal against any belief that assigns a higher probability to $\omega = 1$ than $\omega = 3$.[8] Hence, the high action must be played infinitely often almost surely.

I conclude the 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 $a^\theta, a^{\theta^c}$. Then there exists a p-absorbing SCE under $\theta$.*

It now becomes clear that persisting against a correctly specified model conveys abundant information about $\theta$. Perhaps surprisingly, this alone is powerful enough to guarantee that the model also persists against every other competing model. 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 iff there exists a p-absorbing SCE under $\theta$.*

The seemingly demanding notion of global robustness amounts to the requirement that $\theta$ persists against one arbitrary correctly specified model at some prior and policy. For instance, provided that $\theta$ 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, or any model with arbitrarily complex parameter space. 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 $\theta$, 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.

[8]The optimality follows from the myopic optimality and the symmetry of the setting.

The critical step in proving Theorem 1 shows that the p-absorbing condition implies a stronger condition that I describe in Lemma 3. The condition states that we can find some prior such that, the dogmatic modeler would play the SCE with her belief concentrated around the KL minimizers for all periods with arbitrarily high probability. The key idea is as follows. Fix any $\epsilon > 0$. Define as $\gamma(\pi_0)$ the probability of a dogmatic modeler with model $\theta$ and prior $\pi_0$ only playing actions in $\text{supp}(\sigma)$ and her belief never leaving the $\epsilon$-neighborhood of beliefs over $\Omega^\theta(\sigma)$, denoted by $B_\epsilon(\Delta\Omega^\theta(\sigma))$.[9] When a dogmatic modeler only plays actions in $\text{supp}(\sigma)$, her belief a.s. converges to a limit belief that assigns probability 1 to the KL minimizers $\Omega^\theta(\sigma)$. Therefore, $\gamma(\pi_0) > 0$ for some prior $\pi_0$. By definition, with probability $1-\gamma(\pi_0)$, the agent eventually reaches some posterior $\hat{\pi}_t$ that either induces her to play a non-equilibrium action or leaves the neighborhood $B_\epsilon(\Delta\Omega^\theta(\sigma))$, which means $\gamma(\hat{\pi}_t) = 0$; but this immediately implies there exists some reachable posterior $\tilde{\pi}_t$ at which $\gamma(\tilde{\pi}_t) > \gamma(\pi_0)$, because $\gamma(\pi_0)$ is a weighted average of $\gamma(\pi_t)$ across all reachable posteriors $\pi_t$. Therefore, there exists a sequence of beliefs at which the probability of interest increases to 1.

**Lemma 3.** *If $\sigma$ is a p-absorbing SCE, then given any $\gamma \in (0,1)$ and $\epsilon > 0$, there exists a full-support $\pi_0^\theta$ and an optimal $a^\theta$ under which, with probability higher than $\gamma$, a dogmatic modeler only plays actions in* $\text{supp}(\sigma)$ *and her belief never leaves $B_\epsilon(\Delta\Omega^\theta(\sigma))$ for all periods.*

Going forward, the proof of Theorem 1 shows that the existence of a p-absorbing SCE in fact implies a property even stronger than global robustness: at some prior and policy for $\theta$, with positive probability, the switcher *never* switches regardless of what competing model she entertains. Let $\sigma$ be a p-absorbing SCE under $\theta$. By Lemma 3, there exists a prior $\pi_0^\theta$ and a policy $a^\theta$ under which with probability $\gamma(\pi_0^\theta)$, a dogmatic modeler plays actions in the support of $\sigma$ with her belief staying inside a very small neighborhood of the KL minimizers for all $t$. Since $\sigma$ is self-confirming, it follows that with high probability, the likelihood ratio of any competing model $\theta^c$ to $\theta$ approximates the likelihood ratio of $\theta^c$ to the true DGP. Note the latter likelihood ratio constitutes a supermartingale; by Ville's Maximal Inequality (Ville, 1939), we know that the probability that it exceeds $\alpha > 0$ at least once is bounded above by $1/\alpha$. Therefore, when $\gamma(\pi_0^\theta) > 1 - 1/\alpha$, there exists a positive measure of histories on which the likelihood between $\theta^c$ and $\theta$ never goes beyond $\alpha$. Note that on those histories, the switcher never initiates a switch and her behavior agrees with a dogmatic modeler.

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.[10]

[9] For any set of finite probability distributions $Z$ over sample space $S$, I use $B_\epsilon(Z)$ to denote the set of probability distributions whose minimum distance from any element in $Z$ is smaller than $\epsilon$, i.e. $B_\epsilon(Z) = \{z \in \Delta S : \min_{z' \in Z} d_P(z, z') < \epsilon\}$, where $d_P$ represents the usual Prokhorov metric over $\Delta S$.

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

**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. The proof of Corollary 1 uses a similar technique used by Fudenberg et al. (2021). In particular, Theorem 2 in Fudenberg et al. (2021) implies that if a pure BN-E is uniformly strict, then it is p-absorbing. Corollary 1 further accommodates mixed equilibria by showing that if a mixed BN-E is uniformly quasi-strict, then it is also p-absorbing. Note that the SCE in Example 1 is not uniformly quasi-strict because its support does not include all myopically optimal actions.

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

Example 2 demonstrates how overconfidence in one's ability can give rise to a globally robust misspecified model.

**Example 2** (Overconfidence I)**.** Consider a discrete version of Example 2 in Heidhues et al. (2019). A worker chooses a level of costly effort each period from $\mathcal{A} = \{0, 1, 2, 3\}$ and observes a payoff of $u(a_t, y_t) = y_t - .5a_t^2$. The true DGP determines the output, $y_t = (a_t + b^*)\,\omega^* + \eta_t$, where $b^* = 1$ is his true ability level, $\omega^* = 2$ is an environment fundamental that determines the return to effort and $\eta_t \sim \mathcal{N}(0, 1)$. The efficient effort level is $a^* = 2$. The worker is initially overconfident in that he believes his ability is given by $\hat{b} = 3 > b^*$. He is unsure of the return to effort and needs to learn about it over time. This is reflected from his initial subjective model $\theta$ which treats the return to effort as a parameter to be estimated. According to $\theta$, the output is given by $y_t = \left(a_t + \hat{b}\right)\omega + \eta_t$, where $\omega$ is an element of a finite parameter set $\Omega^\theta = \{1, 2, 3\}$ and $\eta_t \sim \mathcal{N}(0, 1)$.

It can be easily verified that there exists a unique uniformly quasi-strict SCE under $\theta$ and thus model $\theta$ is globally robust. In the SCE, the worker inefficiently chooses $\hat{a} = 1$ and believes in $\hat{\omega} = 1$. Due to overconfidence, he attributes his underperformance to a bad environment and exerts lower effort in response. The interpretation of this result is as follows. Suppose the worker receives a suggestion from his colleagues and starts to consider if he suffers from overconfidence. That is, he now evaluates a competing model $\theta^c$ that is correctly specified about the ability. For example, let $\theta^c$ be correctly specified and predict that $y_t = (a_t + b^*)\,\omega + \eta_t$, where $\Omega^{\theta^c} = \Omega^\theta = \{1, 2\}$. However, 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 $\pi_0^\theta$ that is sufficiently concentrated on $\hat{\omega} = 1$. 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.

An important feature of Theorem 1 is that the characterization of global robustness does not depend on the switching threshold $\alpha$. This stems from the fact that the prior $\pi_0^\theta$ we need to establish global robustness is allowed to vary with $\alpha$. Indeed, the proof of Theorem 1 shows that $\alpha$ is relevant to the choice of the prior. As $\alpha$ decreases, model switching becomes less sticky, so the agent's prior needs to be more entrenched and concentrated around KL minimizers such that the prediction of her initial model stays sufficiently close to the competing model and the probability $\gamma(\pi_0^\theta)$ is sufficiently close to 1. In contrast, when $\alpha$ becomes larger, the agent is more reluctant to switch, allowing her model to persist under more priors.

There arises a natural question as to whether a globally robust model can persist against every competing model under a prior that is not concentrated around the KL minimizers. The answer to this question is *a priori* unclear. Consider a competing model $\theta^c$ that solely consists of the true DGP. If $\theta$ admits a unique SCE $\sigma$ and $\pi_0^\theta(\Omega^\theta(\sigma)) < 1/\alpha$, then conditional on the agent never switching to $\theta^c$, the likelihood ratio between $\theta^c$ and $\theta$ is asymptotically bounded below by $1/\pi_0^\theta(\Omega^\theta(\sigma))$ and thus also by $\alpha$—a switch to $\theta^c$ is destined to occur. While this argument implies the agent must switch at least once, it does not preclude the possibility that she might eventually switch back to $\theta$. In principle, she might draw an outcome that both triggers a switch back and sways her belief towards the KL minimizers, enabling her to stick to $\theta$ thereafter.

Theorem 2 provides a necessary condition that quantifies how concentrated the prior belief must be around the p-absorbing SCE parameters in the model for global persistence. This condition is stated in terms of $\alpha$, and it implies any prior belief that puts positive probability on non-SCE parameters will not lead to global persistence for some low enough switching threshold $\alpha > 1$. An economic implication of this result is that, if a persuader wishes to change a decision maker's view by proposing a competing model, he should do it early—namely, before the decision maker's belief becomes entrenched.

**Theorem 2.** *Denote the set of all p-absorbing SCEs under $\theta$ by $\Sigma^*$. Suppose at a full-support $\pi_0^\theta$ and an optimal $a^\theta$, model $\theta \in \Theta$ persists against every $\theta^c \in \Theta$ at every full-support $\pi^{\theta^c}$ and optimal $a^{\theta^c}$, then $\pi_0^\theta\left(\cup_{\Sigma^*}\Omega^\theta(\sigma)\right) \geq 1/\alpha$.*

The key force behind Theorem 2 is that the Bayes factor rule favors models with a simpler structure. Consider a competing model $\theta^c$ with a parameter set that consists of the KL minimizers in all SCEs, $\cup_{\Sigma^*}\Omega^\theta(\sigma)$, and an additional parameter $\omega^*$ that prescribes the true DGP. Note that $\theta^c$ is correctly specified with the addition of $\omega^*$. If we ignore the newly added parameter $\omega^*$, this competing model is simpler than $\theta$ as it prescribes a smaller set of outcome distribution profiles. Let the prior for the competing model assigns a negligible probability to $\omega^*$ and, to the parameters it shares with $\theta$, probabilities proportional to those assigned by $\pi_0^\theta$. It follows that the likelihood ratio between $\theta^c$ and $\theta$ is asymptotically bounded below by the posterior-prior ratio over the KL minimizers, i.e. $\pi_t^\theta(\cup_{\Sigma^*}\Omega^\theta(\sigma))/\pi_0^\theta(\cup_{\Sigma^*}\Omega^\theta(\sigma))$. Lemma 2 tells us that if $\theta$ persists, then the belief eventually becomes concentrated on the KL minimizers, but this immediately implies that the likelihood ratio between $\theta$ and $\theta$ must rise above $\alpha$ if $\pi_0^\theta\left(\cup_{\Sigma^*}\Omega^\theta(\sigma)\right) < 1/\alpha$, a contradiction.

## 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. Under those circumstances, global robustness may not be the best criterion. A natural question is then whether a model persists against every possible *neighbor* model. When the answer is no, then such a model cannot be adopted forever even when the agent evaluates an only slightly different model. By contrast, if the answer is yes, then this model can persist whenever the switcher only takes small steps. Following this line of thought, we can develop a weaker robustness property, *local robustness.*

First, I formalize what qualifies as a neighbor model. In this and next subsection, 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 non-parametric, and therefore, *unconstrained.* In the next subsection, I introduce a parametric approach that places additional 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,

$$d\left(Q, Q'\right) \coloneqq \max_{a\in\mathcal{A}} d_P\left(Q_a, Q'_a\right), \tag{8}$$

where

$$d_P\left(Q_a, Q'_a\right) = \inf\left\{\epsilon > 0 | Q_a\left(Y\right) \leq Q'_a\left(B_\epsilon\left(Y\right)\right) + \epsilon \text{ for all } Y \subseteq \mathcal{Y}\right\}. \tag{9}$$

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 $\theta$, and let $\mathcal{Q}^\theta \coloneqq$

$\left\{Q^{\theta,\omega}\right\}_{\omega\in\Omega^\theta}$ denote all distributions in the support of $\theta$. Define the $\epsilon$-neighborhood of $\theta$ as

$$N_\epsilon(\theta) \coloneqq \left\{\theta' \in \Theta : d_H\left(\mathcal{Q}^\theta, \mathcal{Q}^{\theta'}\right) \le \epsilon\right\}, \tag{10}$$

where

$$d_H\left(\mathcal{Q}^\theta, \mathcal{Q}^{\theta'}\right) \coloneqq \max\left\{\max_{\omega\in\Omega^\theta}\min_{\omega'\in\Omega^{\theta'}} d\left(Q^{\theta,\omega}, Q^{\theta',\omega'}\right), \max_{\omega'\in\Omega^{\theta'}}\min_{\omega\in\Omega^\theta} d\left(Q^{\theta,\omega}, Q^{\theta',\omega'}\right)\right\} \tag{11}$$

denotes the Hausdorff metric. Besides restricting the distance between models, I argue that a notion of local robustness should also put limit on the distance between the priors for the initial model and the competing neighbor model. To see this, imagine an agent, who originally believes that with probability .99 vaccines do not work, suddenly evaluates a competing model with a prior assigning the same probability .99 to the opposite scenario that vaccines actually work—this requires a bit of a leap of faith. To measure the distance over beliefs, note that every prior $\pi$ over a parameter space $\Omega^\theta$ corresponds to a belief over $(\Delta\mathcal{Y})^\mathcal{A}$; denote this mapping as $\Gamma^\theta : \Delta\Omega^\theta \to \Delta(\Delta\mathcal{Y})^\mathcal{A}$. Denote the set of beliefs for model $\theta^c$ that are within $\epsilon$ distance from $\pi^\theta$ by

$$D_\epsilon(\pi^\theta;\theta^c) \coloneqq \{\pi^{\theta^c} \in \Delta\Omega^{\theta^c} : d_P\left(\Gamma^\theta(\pi^\theta), \Gamma^{\theta^c}(\pi^{\theta^c})\right) \le \epsilon\}. \tag{12}$$

I now define unconstrained local robustness.

**Definition 5.** A model $\theta$ is *unconstrained locally robust* if there exists some $\epsilon > 0$, a full-support $\pi_0^\theta$, and an optimal $a^\theta$ under which $\theta$ persists against every $\theta^c \in N_\epsilon(\theta)$ at every full-support $\pi_0^{\theta^c} \in D_\epsilon(\pi_0^\theta;\theta^c)$ and optimal $a^{\theta^c}$.

The definition of local robustness seems much weaker than global robustness. It only requires the model to be able to persist against neighbor models at nearby beliefs. When $\theta$ is misspecified, a sufficiently close model is necessarily also misspecified. This prevents us from invoking Lemmas 1 and 2 and inferring that a p-absorbing SCE must exist. However, perhaps surprisingly, Theorem 3 shows that these two notions of robustness point to the same set of subjective models.

**Theorem 3.** *Unconstrained local robustness is equivalent to global robustness.*

The idea of Theorem 3 is quite simple. Given any unconstrained locally robust model $\theta$, we could construct a neighbor competing model $\theta^c$ with the same parameter set and the same prior but potentially different outcome distributions, such that each parameter prescribes a convex combination of the true DGP and the corresponding distribution under $\theta$. Then the likelihood ratio of $\theta^c$ to $\theta$ would also be a linear combination of 1 and the likelihood ratio of the true DGP to $\theta$. Hence, $\theta$ persists against $\theta^c$ only if $\theta$ also persists against the true DGP, which makes it globally robust. Theorem 3 also provides a new perspective for understanding global robustness. If model $\theta$ 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 unconstrained local robustness introduced in the previous section has not restricted the nature of models that are being compared by the switcher. I now take a parametric approach by restricting attention to models sharing the same parametric structure. Take any profile of action-dependent outcome distributions, $\{q(\cdot|a,\omega)\}_{a\in\mathcal{A},\omega\in\Omega^q}$, that are uniformly continuous over $\Omega^q$ for all $a \in \mathcal{A}$, where the parameter set $\Omega^q$ is any subset of a Euclidean space (potentially infinite and unbounded). Define its affiliated family of models as follows.

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

$$\Theta^q := \{\theta \in \Theta : q^\theta(\cdot|a,\omega) \equiv q(\cdot|a,\omega) \text{ for all } \omega \in \Omega^\theta \subseteq \Omega^q \text{ and all } a \in \mathcal{A}\}. \tag{13}$$

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

$$N^q_\epsilon(\theta) := \left\{\theta' \in \Theta^q : d_H\left(\Omega^\theta, \Omega^{\theta'}\right) \le \epsilon\right\}, \tag{14}$$

where $d_H\left(\Omega^\theta, \Omega^{\theta'}\right) = \max\left\{\max_{\omega\in\Omega^\theta}\min_{\omega'\in\Omega^{\theta'}}\|\omega-\omega'\|, \max_{\omega'\in\Omega^{\theta'}}\min_{\omega\in\Omega^\theta}\|\omega-\omega'\|\right\}$. The distance measure for priors is also significantly simpler. Denote the $\epsilon$-neighborhood of any $\pi \in \Delta\Omega^q$ by

$$D^q_\epsilon(\pi) := \{\pi' \in \Delta\Omega^q : d_P(\pi,\pi') \le \epsilon.\} \tag{15}$$

Now we can define constrained local robustness in a similar fashion.

**Definition 7.** A model $\theta \in \Theta^q$ is *q-constrained locally robust* if there exists some $\epsilon > 0$, a full-support $\pi^\theta_0$, and an optimal $a^\theta$ under which $\theta$ persists against every $\theta^c \in N^q_\epsilon(\theta)$ at every full-support $\pi^{\theta^c}_0 \in D^q_\epsilon(\pi^\theta_0)$ and optimal $a^{\theta^c}$.

The notion of $q$-constrained local robustness requires a model to be able to persist against all nearby models within the $q$-family under nearby priors. By embedding a model into different parametric families, the switcher is not limited to merely considering an expansion or a downsizing of the parameter set of her initial model. Below I present an example to illustrate the flexibility of the approach.

**Example 3** (Overconfidence II)**.** I now extend Example 2 to a teamwork setting. Similarly, the worker chooses how much effort to exert from the set $\mathcal{A} = \{0, 1, 2, 3\}$ and obtains a flow payoff $u(a_t, y_t) = y_t - .5a_t^2$. At the end of each period, he observes the joint output of his team $y_t = (a_t + b^*)\omega^* + \eta_t$, where $b^* = 1$ is the worker's true ability and $\eta_t \sim \mathcal{N}(0, 1)$. In contrast to Example 2, let $\omega^* = 3$ be the ability of his teammate and note that the worker is much less capable than his teammate. In addition to the output, the worker receives periodic performance reviews from his supervisor, taking the form of $s_t = b^{*2} + \xi_t$ with $\xi_t \sim \mathcal{N}(0, 2)$.

The worker has an initial subjective model $\theta$ with $\Omega^\theta = \{1, 2, 3\}$ and presumes that his ability is given by $\hat{b} = 2 > b^*$. Since the performance review provides access to an infinite sequence of unbiased signals about $b^*$, his model $\theta$ admits no self-confirming equilibrium and, by Theorem 1, is not globally robust.

By specifying a parametric family for $\theta$, we essentially define the decision maker's consideration set of alternative neighbor models. For example, the worker may want to explore perturbations about his ability while maintaining the assumption that his ability is no lower than his counterpart's. This may arise when worker suffers from the "better-than-average" effect (Langer and Roth, 1975; Svenson, 1981; Benoît, Dubra, and Moore, 2015). To capture this, let $q(\cdot|a, b, \omega)$ be the joint distribution of $y_t$ and $s_t$ when the worker takes action $a$ and the ability levels are given by $(b, \omega)$, and define $\Omega^q = \{(b, \omega) \in \mathbb{R}^2 : b \geq \omega\}$. Although the parameter space $\Omega^\theta$ is one-dimensional, we can equivalently write it as $\overline{\Omega}^\theta = \{\hat{b}\} \times \Omega^\theta$. With this augmented parameter space, model $\theta$ belongs to the $q$-family. Alternatively, we could remove the restriction and specify $\Omega^q = \mathbb{R}^2$. In this case, $q$-constrained local robustness requires the worker's model to persist against all perturbations about his and his teammate's ability levels.

### 4.3.1 Sufficient Conditions

In this subsection, I provide sufficient conditions for $q$-constrained local robustness. Several complicating issues arise as a result of the constraints. First of all, as in the unconstrained local robustness case, models sufficiently close to a misspecified model must also be misspecified, preventing us from using Lemma 2. Due to the additional constraints over the parametric structure, it is now infeasible to perturb the predictions of model $\theta$ 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.

Similar to the idea of Theorem 1, model $\theta$ may be $q$-constrained locally robust if we can verify that $\theta$ admits a p-absorbing BN-E $\sigma$ such that when the agent plays the actions prescribed by $\sigma$, with positive probability, the agent never finds a neighbor model more attractive. This would be satisfied if $\sigma$ is pure and the likelihood ratio between any parameter in $\Omega^q$ that neighbors $\Omega^\theta$ and any KL minimizer in $\Omega^\theta(\sigma)$ remains under $\alpha$ forever with positive probability.[11] It comes for free when $\sigma$ is self-confirming as the aforementioned likelihood ratio is a supermartingale, but requires additional assumptions otherwise. Frick et al. (2021a) develops an order of prediction accuracy that partially restores the martingale argument. I now define a property based on the prediction accuracy order and generalize their definition to accommodate a non-singleton $\Omega^\theta(\sigma)$.

**Definition 8.** Model $\theta \in \Theta^q$ is *locally dominant* at $\sigma$ w.r.t. $\Omega^q$ if there exists $\epsilon > 0$ and $d > 0$ such that for all $\omega \in \Omega^\theta(\sigma)$, all $\omega' \in \Omega^q \cap B_\epsilon(\Omega^\theta(\sigma))$, and all $a \in \operatorname{supp}(\sigma)$, we have $\mathbb{E}\left(\frac{q(\cdot|a,\omega')}{q(\cdot|a,\omega)}\right)^d \leq 1$.

Local dominance of model $\theta$ at $\sigma$ implies that there exists $d > 0$ such that for any KL-minimizer $\omega$ in $\Omega^\theta(\sigma)$, when actions in the support of $\sigma$ are played, the $d$-th power of the likelihood ratio between any neighbor parameter in $\Omega^q$ and $\omega$ is a supermartingale. I now state a sufficient condition for $q$-constrained local robustness.

**Theorem 4.** *A model $\theta \in \Theta^q$ is $q$-constrained locally robust if $\theta$ admits a pure p-absorbing BN-E $\sigma$ at which $\theta$ is locally dominant w.r.t. $\Omega^q$.*

The proof idea of Theorem 4 is as follows. When a competing model $\theta^c$ is sufficiently close to the initial model $\theta$, we can decompose its parameter space $\Omega^{\theta^c}$ into $|\Omega^\theta|$ disjoint subsets, each being contained in a small neighborhood of some parameter in $\Omega^\theta$. For each $\omega \in \Omega^\theta$, denote its corresponding neighbor parameters in $\Omega^{\theta^c}$ by $\iota(\omega)$. Then by definition, we have $\Omega^{\theta^c} = \cup_{\omega \in \Omega^\theta} \iota(\omega)$. Suppose for simplicity that there is a unique KL minimizer at the BN-E $\sigma$, $\omega^* \in \Omega^\theta(\sigma)$. As an implication of local dominance and Ville's maximal inequality

[11] Note that additional conditions are needed if $\sigma$ is not pure. Suppose $\sigma$ is a p-absorbing mixed BN-E. Unless $\sigma$ is self-confirming, when the agent only plays actions in the support of $\sigma$, the parameters that empirically best fit the observed data can change with the empirical action frequency and are not necessarily be given by $\Omega^\theta(\sigma)$. Hence, the likelihood ratio between a parameter in $\Omega^q$ and $\theta$ cannot be approximated by the likelihood ratio between that parameter and a KL minimizer in $\Omega^\theta(\sigma)$. In a more generalized version of Theorem 4, $q$-constrained local robustness is implied by the existence of a p-absorbing mixed BN-E $\sigma$ such that $\Omega^\theta(\sigma') = \Omega^\theta(\sigma)$ for each $\sigma' \in \Delta \operatorname{supp}(\sigma)$ and $\theta$ is locally dominant at $\sigma$ w.r.t. $\Omega^q$.

for supermartingales, for each $\omega \in \Omega^\theta$, the likelihood ratio

$$\frac{\sum_{\omega' \in \iota(\omega)} \pi_0^{\theta_c}(\omega') \prod_{\tau=1}^t q(y_\tau | a_\tau, \omega')}{\prod_{\tau=1}^t q(y_\tau | a_\tau, \omega^*)}$$

remains under $\kappa \cdot \pi_0^{\theta^c}(\iota(\omega))$ for any constant $\kappa > 1$ with positive probability if $a_t \in \text{supp}(\sigma)$ for all $t$, and the probability approaches 1 as $\kappa$ increases to infinity. I then show that, when the prior $\pi_0^\theta$ is sufficiently concentrated at the KL minimizer $\omega^*$ and the prior $\pi_0^{\theta^c}$ is sufficiently close to $\pi_0^\theta$, the sum of such likelihood ratios across all $\omega \in \Omega^\theta$ remains under $\alpha$ forever with positive probability. From here, we can apply the argument for Theorem 1 and show $\theta$ persists against $\theta^c$.

Analogous to Corollary 2, the following corollary provides a sufficient condition that can be verified from the primitives.

**Corollary 3.** *A model $\theta \in \Theta^q$ is q-constrained locally robust if $\theta$ admits a uniformly strict BN-E $\sigma$ at which $\theta$ is locally dominant w.r.t. $\Omega^q$.*

Let's revisit Example 3 to illustrate how to apply the results.

**Example 3, cont.** Suppose the worker indeed maintains the assumption that he has weakly higher ability than his teammate and evaluates perturbations within $\Omega^q = \{(b, \omega) \in \mathbb{R}^2 : b \geq \omega\}$. Despite that the worker's initial model $\theta$ is not globally robust, I now show that it is $q$-constrained locally robust. Note that $\theta$ gives rise to a uniformly strict Berk-Nash equilibrium in which the worker chooses $\hat{a} = 2$ and believes in $\hat{\omega} = 2$. Model $\theta$ does not yield a perfect match to the outcome distributions, as the predicted average output is 8 and the predicted average review is 4, different from the actual average output of 9 and the actual average review of 1. However, for any $(b, \omega)$ such that $b \geq \omega$, we have

$$\begin{aligned}
\mathbb{E}_{y,s}\left(\frac{q(y, s|\hat{a}, b, \omega)}{q(y, s|\hat{a}, \hat{b}, \hat{\omega})}\right) &= \exp\left([(2+b)\omega - (2+\hat{b})\hat{\omega}][(2+b^*)\omega^* - (2+\hat{b})\hat{\omega}] + \frac{1}{2}(b^2 - \hat{b}^2)(b^{*2} - \hat{b}^2)\right) \\
&= \exp\left((2+b)\omega - \frac{3}{2}b^2 - 2\right) \leq \exp\left(-\frac{1}{2}(b-2)^2\right) \leq 1.
\end{aligned}$$

By Theorem 4, model $\theta$ is $q$-constrained locally robust. Although the worker now deems it possible that his ability may be lower than expected and a lower ability would indeed better explain the performance reviews, this possibility also forces him to lower his assessment about his teammate's ability due to the "better-than-average" bias, implying a worse fit to the output realizations. When the worker sustains an intermediate level of overconfidence and believes in $\hat{b} = 2$, the latter force dominates and his model persists against neighbor models within the "better-than-average" class.

#### 4.3.2 Necessary Conditions

Restricting the set of competing models and priors also poses a challenge to finding necessary conditions for local robustness. By Lemma 2, when the competing model is correctly specified and the initial model is eventually adopted, the agent's action must converge to the support of a Berk-Nash equilibrium. This greatly simplifies the characterization of persistence because we could simply compare $\theta$ with the competing model at the BN-E. This convenient convergence property is lost when we shift our focus to misspecified competing models. However, as I show in Theorem 5, suppose the action frequency of a dogmatic-modeler with $\Theta^{\dagger} = \{\theta\}$ indeed a.s. converges, then it is necessary for a $q$-constrained locally robust model to admit at least one BN-E $\sigma$ at which its KL minimizers $\Omega^{\theta}(\sigma)$ fit the data better than neighbor parameters in the sense that they locally minimize the KL divergence within $\Omega^{q}$.[12]

**Definition 9.** Model $\theta$ is *locally KL-minimizing* at $\sigma$ w.r.t. $\Omega^{p}$ if there exists $\epsilon > 0$ such that for all $\omega \in \Omega^{\theta}(\sigma)$ and $\omega' \in \Omega^{q} \cap B_{\epsilon}\left(\Omega^{\theta}(\sigma)\right)$, we have $\sum_{\mathcal{A}} \sigma(a) \mathbb{E}\left(\ln \frac{q(\cdot|a,\omega')}{q(\cdot|a,\omega)}\right) \leq 0$.

Local KL-minimization is weaker than local dominance in two aspects. First, local KL-minimization only compares nearby parameters at the mixed action $\sigma$ while local dominance makes a comparison at each action in the support of $\sigma$. Second, fixing an action $a$, as Frick et al. (2021a) point out, if $\mathbb{E}\left(\frac{q(\cdot|a,\omega')}{q(\cdot|a,\omega)}\right)^{d} \leq 1$ for any $d > 0$, then we immediately have $\mathbb{E}\left(\ln \frac{q(\cdot|a,\omega')}{q(\cdot|a,\omega)}\right) < 0$.[13] While local dominance enables us to invoke maximal inequalities and show the likelihood ratio remains under $\alpha$ forever with positive probability, the violation of local KL-minimization is all we need to establish that the likelihood ratio almost surely exceeds $\alpha$.

**Theorem 5.** *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 $\theta \in \Theta^{q}$ is $q$-constrained locally robust only if it admits a BN-E $\sigma$ at which $\theta$ is locally KL-minimizing w.r.t. $\Omega^{q}$.*

Esponda et al. (2019) establish global almost-sure convergence of a dogmatic modeler's action frequency to a BN-E if it is "globally attracting" and the dogmatic modeler is myopic ($\delta = 0$), where 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 couple

[12] 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 $\sigma_t(a) = \frac{1}{t}\sum_{\tau=1}^{t} 1_{\{a_t = a\}}$.

[13] Moreover, the fact that $\mathbb{E}\left(\ln \frac{q(\cdot|a,\omega')}{q(\cdot|a,\omega)}\right) < 0$ holds in a small neighborhood of $\omega$ does not imply there exists $d > 0$ such that $\mathbb{E}\left(\frac{q(\cdot|a,\omega')}{q(\cdot|a,\omega)}\right)^{d} \leq 1$ for all $\omega'$ in the neighborhood.

of 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 5 provides a simple criterion to determine if some subjective model is constrained locally robust.

The proof idea of Theorem 5 is as follows. Suppose none of the (potentially infinitely many) Berk-Nash equilibria under $\theta$ satisfy the local KL-minimization property w.r.t. $\Omega^q$. Then for each equilibrium $\sigma$, there exists a parameter $\omega' \in \Omega^q \backslash \Omega^\theta$ that is close to some $\omega \in \Omega^\theta$ and an open neighborhood of $\sigma$, $B_\epsilon(\sigma) \subseteq \Delta\mathcal{A}$, such that at any strategy $\sigma' \in B_\epsilon(\sigma)$, the parameter $\omega'$ yields a strictly lower KL divergence than the lowest possible KL divergence under $\theta$. Since the set of Berk-Nash equilibria is compact (Lemma 8), the Heine-Borel theorem implies that there exists a finite set of parameters $\Omega' \subseteq \Omega^q \backslash \Omega^\theta$ such that for each equilibrium $\sigma$, there exists an $\omega' \in \Omega'$ that satisfies this property. Consider the competing model $\theta^c \in \Theta^q$ with a larger parameter space $\Omega^{\theta^c} = \Omega^\theta \cup \Omega'$. Suppose the model choice of the agent converges to $\theta$ with positive probability. On the set of paths where she eventually adopts $\theta$ forever, she behaves identically to a dogmatic modeler after the final switch to $\theta$, and thus, by assumption, her action frequency converges to a Berk-Nash equilibrium. Under such convergence, $\theta^c$ strictly outperforms the fit of $\theta$ in the long term. Therefore, $\lambda_t^{\theta^c}$ eventually exceeds $\alpha$ and the switcher adopts $\theta^c$. This is a contradiction. Theorem 5 follows.

It is challenging to identify necessary conditions without global convergence of a dogmatic modeler's behavior. Nevertheless, I show in the Appendix that the local KL-minimization condition is still necessary when $\mathcal{A}$ is binary and the agent is myopic (Theorem 7). I conclude this section by returning to Example 3.

**Example 3, cont.** Suppose the worker is open to evaluate all competing models without the restriction that his ability is weakly higher than his teammate's. A $q$-constrained locally robust model must involve a correct self-perception. Take any $\sigma \in \Delta\mathcal{A}$, we have

$$\sum_{\mathcal{A}} \sigma(a)\mathbb{E}\left(\ln q(y, s|a, b, \omega)\right) = \sum_{\mathcal{A}} \sigma(a)\frac{1}{2}[(a+b)\omega - (a+b^*)\omega^*]^2 + \frac{1}{4}(b^2 - b^{*2}),$$

which is locally minimized only at $(b^*, \omega^*)$. Since this holds for any $\sigma$, by a similar argument as in the proof of Theorem 5, it can be shown that a model $\theta$ is $q$-constrained locally robust only if $\hat{b} = b^*$ and $\omega^* \in \Omega^\theta$.

# 5 Multiple Competing Models

The framework described in Section 2 could be easily extended to accommodate multiple competing models. Let $\Theta^c \subseteq \Theta$ denote a finite subset of competing models and $\Theta^\dagger = \Theta^c \cup \theta$

denote the set of all subjective models the agent entertains. Replacing $\theta^c$ by $\Theta^c$ in all previous definitions, we obtain a framework that allows for switching between more than two models. In particular, the agent compares her current model against all alternatives and switches to the one with the highest likelihood ratio if it exceeds $\alpha$. Suppose the agent entertains at most $K \geq 1$ competing models, then global robustness would require $\theta$ to persist against every $\Theta^c \subseteq \Theta$ of size no larger than $K$ at all priors and policies assigned to models in $\Theta^c$; $q$-constrained local robustness would require $\theta$ to persists against every set of neighbor models $\Theta^c$ of size no larger than $K$ within the $q$-family at all neighbor priors and policies for $\Theta^c$.[14]

The first insight that emerges from this extension is that the evaluation of multiple competing models is prone to overfitting. In particular, when the switching threshold $\alpha$ is not adjusted as the number of competing models $K$ becomes larger, even the true DGP may fail to be globally robust. Relatedly, Schwartzstein and Sunderam (2019) find in a static setting, a decision maker switches from the true DGP to a competing model when a persuader is allowed to propose one after the data is realized. In contrast, I show that the persuader can achieve the same goal even if he has to propose before the outcomes are drawn, provided that he can present multiple competing models. Additionally, as illustrated by Example 4 below, depending on the circumstances, not only must the switcher switch at least once to the competing models, due to the stickiness of model switching, she may also eventually settle down with one of them despite that the true DGP fits the data perfectly on average.

Let $M$ be the smallest integer such that $M > \alpha + 1$. I now construct an example in which there are $M$ competing models and the true DGP does not persist at any priors and policies.

**Example 4** (Overfitting)**.** The agent has two actions $\mathcal{A} = \{a', a''\}$. The true DGP is a uniform distribution over $M$ outcomes, $\mathcal{Y} = \{1, ..., M\}$, regardless of the action. The agent incurs a loss of $-M$ when drawing the outcome $y = 1$ but a payoff of 0 if another outcome is realized. The agent pays an additional cost $c > 0$ for playing $a'$ and no cost if she plays $a''$. Hence, if the true DGP is known, the agent would play $a''$ to avoid the cost.

The agent has the true DGP as her initial model $\theta$. According to the prediction of $\theta$, the agent optimally plays $a''$ in the first period. Suppose the agent evaluates $M$ competing models that I describe below. Each model $\theta^n \in \{\theta^1, ..., \theta^M\}$ has a single parameter $\omega^n$, with $q^{\theta^n}(\cdot|a', \omega^n)$ corresponding to a uniform distribution over $\mathcal{Y}$. When $n \neq 1$, $q^{\theta^n}(\cdot|a'', \omega^n)$ is

[14] If $\Theta^c$ is not a singleton, then persistence against $\Theta^c$ is not equivalent to persistence against each model in $\Theta^c$, and neither implies the other. See Appendix C for examples.

given by

$$q^{\theta^n}(y|a'',\omega^n) = \begin{cases} 1-\frac{1}{M}-(M-1)\eta & \text{if } y=n, \\ \frac{1}{M}+\eta & \text{if } y=1, \\ \eta & \text{if } y\in\mathcal{Y}\setminus\{1,n\}, \end{cases}$$

where $\eta$ is a small positive constant. When $n=1$, $q^{\theta^n}(\cdot|a'',\omega^n)$ is analogously given by

$$q^{\theta^1}(y|a'',\omega^n) = \begin{cases} 1-(M-1)\eta & \text{if } y=1, \\ \eta & \text{if } y\in\mathcal{Y}\setminus\{1\}. \end{cases}$$

Model $\theta^n$ predicts that when $a''$ is played, the outcome $n$ is realized with probability near 1. Since there is one such model for every potential outcome, the agent must switch to one of these competing models after the first period. In particular, suppose the realized outcome is $n$, then a switch to model $\theta^n$ is triggered immediately when $\eta$ is sufficiently small such that

$$\frac{l_1^{\theta^n}}{l_1^{\theta}} \geq \frac{1-\frac{1}{M}-(M-1)\eta}{\frac{1}{M}} > \alpha.$$

Since each competing model assigns a probability larger than $1/M$ to the outcome 1, once the switch happens, the agent finds it optimal to play $a'$ instead to avoid the loss associated with outcome 1 if $c$ is sufficiently small. However, as all models have the same correct prediction under $a'$, the likelihood ratios remain unchanged thereafter. Hence, despite initially having the true model, the agent becomes permanently trapped with a wrong model and inefficient play.

In the above example, while the true DGP perfectly matches the outcome distribution in the long term, overfitting arises in the short term and prompts an early switch to the competing models. The more competing models the agent evaluates, the more likely such a switch occurs. One intuitive solution is to make switching stickier so that agent is less responsive to early draws. Indeed, Theorem 6 shows that if we bound $\alpha$ by $K$ from below, the scope of global robustness does not change at all. Theorems 1 and 3 still fully characterize the set of globally robust and unconstrained robust models. A larger bound may be needed for $q$-constrained locally robustness, but when $\alpha$ is large enough, Theorem 4 also generalizes to this environment.

**Theorem 6.** *Suppose the agent evaluates at most $K$ competing models.*

1. *When $\alpha > K$, model $\theta \in \Theta$ is globally robust (or unconstrained locally robust) iff there exists a p-absorbing SCE under $\theta$.*

2. *When $\alpha > K^{1/d}$ for some $d > 0$, model $\theta \in \Theta^q$ is q-constrained locally robust if $\theta$ admits a pure p-absorbing BN-E at which $\theta$ is locally dominant w.r.t. $\Omega^q$ and the local dominance condition in Definition 8 holds at d.*

This extension also highlights the advantage of the Bayes factor rule when it comes to guarding against overfitting. A competing model potentially consists of multiple outcome distribution profiles, each prescribed by a particular parameter value. A prior is assigned to the parameters, which the agent uses as weights when computing the likelihood ratios between two models. A model with more parameters (outcome distributions) also tends to have a more diffuse prior, imposing a penalty for including two many possibilities. This contrasts to the case of multiple competing models, where multiple likelihood ratios are calculated and tracked, and the agent makes a switch as long as the maximum likelihood ratio exceeds the switching threshold, resulting in overfitting.

# 6 Applications

I present two applications to demonstrate how results in this paper uncover new insights about which forms of model misspecification persist and which do not. In a financial investment setting, extreme optimism is shown to be not robust while extreme pessimism is robust. The second application examines an economic environment where an incorrect self-perception distorts a worker's assessment about an outside fundamental and shows that overconfidence is more likely to be robust than underconfidence because they give rise to Berk-Nash equilibria with distinct stability properties.

## 6.1 Pessimism and Optimism

This application compares pessimism and optimism in a setting of financial investing. I show that with mild levels of pessimism or optimism, the investor's subjective model is globally robust; extreme pessimism continues to be globally robust, whereas extreme optimism is no longer constrained locally robust. This follows from the observation that extreme pessimism drives the investor out of the stock market but extreme optimism motivates the investor to actively invest and find out the true market condition.

Consider an investor who has access to $N$ risky assets, indexed by $\{1, 2, ..., N\}$, and one safe asset, indexed by $N+1$. For simplicity, I assume that he invests in only one asset in each period, with his action $a_t \in \mathcal{A} = \{1, 2, ..., N+1\}$ specifying his asset choice. The investor has utility function $u(a_t, y_t) = y_t$, where $y_t$ is the investor's return.[15] The safe asset always

[15] Risk neutrality is only assumed for simplicity of exposition. One could replace $u$ with an arbitrary

returns a known constant $G$. In contrast, the risky assets have random returns with unknown and potentially different means. By investing in the risky asset $n$, the investor's return $y_t$ is the sum of a zero-mean random noise with a known normal distribution and a deterministic mean function $g(b^*, \omega_n^*)$, where $b^* \in [\underline{b}, \overline{b}]$ represents a market factor common to all assets (e.g. the GDP growth of the economy), $\omega_n^* \in [\underline{\omega}, \overline{\omega}]$ represents an asset-specific factor, and $g$ is strictly increasing and continuous in both arguments. Let $g^* := \max_{1 \leq n \leq N} g(b^*, \omega_n^*)$ denote the maximum average return of the risky assets. I make the following assumption about the asset returns.

**Assumption 3.** *$g(\overline{b}, \underline{\omega}) > G$, $g(\underline{b}, \overline{\omega}) < G$, and $g^* > G$.*

The first two inequalities ensure that an extremely favorable (unfavorable) market condition renders all risky assets more (less) profitable than the safe asset. The third inequality implies the true state of the world is such that at least one of the risky assets is more profitable than the safe asset.

I focus on misperceptions about the market condition $b^*$ and models that assign probability 1 to $\hat{b} \in [\underline{b}, \overline{b}]$. If $\hat{b} > b^*$ ($\hat{b} < b^*$), then the investor is optimistic (pessimistic) about the market condition. At the same time, the investor is unsure about the asset-specific variations and treats them as parameters to be estimated. Let the $q$-family be the set of models that share the same structure as the initial model but are associated with potentially different parameter spaces and perceptions about the market condition.[16] Consider a finite set of possible misperceptions $\mathcal{B} \subseteq [\underline{b}, \overline{b}]$. To avoid trivial cases of non-robustness, I assume that for all $b' \in \mathcal{B}$ and all $n$, if there exists some $\omega_n' \in [\underline{\omega}, \overline{\omega}]$ such that $g(b', \omega_n') = g(b^*, \omega_n^*)$, then the investor assigns positive probability to the parameter value $\omega_n'$. We have the following result.

**Proposition 1.** *There exist $\underline{\beta}, \overline{\beta} \in [\underline{b}, \overline{b}]$ such that $\underline{\beta} < b^* < \overline{\beta}$, $g(\underline{\beta}, \overline{\omega}) = G$, and $g(\overline{\beta}, \underline{\omega}) = g^*$. For all $\hat{b} \in \mathcal{B}$, the following statements hold:*

1. *Models with mild optimism or mild pessimism such that $\hat{b} \in [\underline{\beta}, \overline{\beta}]$ are globally robust.*

2. *Models with extreme pessimism such that $\hat{b} < \underline{\beta}$ are globally robust; an investor with one such initial model invests in the safe asset forever.*

3. *Models with extreme optimism such that $\hat{b} > \overline{\beta}$ are not $q$-constrained locally robust. Given any model $\theta$ with $\hat{b} > \overline{\beta}$ and any $\tilde{b} \in \mathcal{B} \cap [b^*, \hat{b})$, there exists a model $\theta^c \in \Theta^q$ associated with $\tilde{b}$ such that model $\theta$ does not persist against $\theta^c$.*

strictly increasing function and obtain the same results with a slight modification of Assumption 3.

[16] A formal mathematical construction is contained in the Appendix.

When the investor has a very low belief about the market condition such that the highest possible payoff from the risky assets is lower than $G$, he always invests in the safe option as it is strictly optimal. He then stops receiving signals about the risky returns, which constitutes a uniformly strict SCE and allows the misspecification to persist. By comparison, an investor with a very high belief finds that the lowest possible payoff from investing in the risky assets is higher than $G$, but actively investing in turn leads him to realize the discrepancy between the expected high returns and the actual returns; hence, no SCE exists. Further, I show that his model is not $q$-constrained locally robust since a competing model with a slightly lower (thus more correct) perception about the market factor will appear more compelling. Finally, mild levels of optimism and pessimism are also globally robust. To see that, note that investing in the most profitable risky assets is always a quasi-strict SCE supported by a mildly wrong belief about the asset-specific factors for assets purchased in equilibrium.[17]

## 6.2 Overconfidence and Underconfidence

This application adapts the setting of Heidhues et al. (2018) to the model-switching environment and compares overconfidence and underconfidence. An incorrect self-perception distorts a worker's belief about his teammate's ability. Higher beliefs over the teammate's ability lead to a monotonic change in the optimal response and, depending on the direction of the biases, may result in even higher beliefs or end up inducing lower beliefs. I demonstrate that overconfidence, featuring positive reinforcement of beliefs, is globally robust, while underconfidence, featuring negative reinforcement of beliefs, is globally robust only on a sequence of unconnected intervals.

This result breaks the symmetry between overconfidence and unconfidence and provides a novel mechanism for why one bias might be expected to be more prevalent than the other. A plethora of evidence in psychology and economics suggests that overconfidence is more prevalent than underconfidence, and many hold the view that this is because agents derive ego utility from holding overconfident beliefs about their own positive traits (Brunnermeier and Parker, 2005; Köszegi, 2006; Oster, Shoulson, and Dorsey, 2013). In contrast, the result in this section provides a reason rooted in the misspecified learning environment itself: overconfidence has better stability properties than underconfidence when the agent can switch models.

[17] The same framework also directly applies to other multi-armed bandit problems. For example, one can imagine a worker who chooses which project to undertake and has a wrong belief about his ability while remaining unsure about the return of each individual project. Another example may be a pharmaceutical firm that has to decide between different development plans for a drug but has an inaccurate assessment about its potential efficacy. Similar to the financial investing problem, extreme misspecification is globally robust if they induce the play of a safe outside option with a known payoff.

As in Example 2, a worker chooses between different effort levels $a_t$ from a finite set $\mathcal{A}$ in each period. Similar to the previous application, the agent has utility function $u(a_t, y_t) = y_t$, where $y_t$ is the output of his work. The output takes the form of $y_t = g(a_t, b^*, \omega^*) + \eta_t$, where $\eta_t$ follows a known mean-zero normal distribution $f$, $b^* \in [\underline{b}, \overline{b}]$ is the agent's ability, and $\omega^* \in [\underline{\omega}, \overline{\omega}]$ is a measure of his teammate's ability. The function $g$ satisfies the following assumption.

**Assumption 4.** *Function $g$ is twice continuously differentiable and strictly increasing in both $b$ and $\omega$. In addition,*

1. *for all $a \in \mathcal{A}$ and all $b \in [\underline{b}, \overline{b}]$, there exists $\omega \in [\underline{\omega}, \overline{\omega}]$ such that $g(a, b, \omega) = g(a, b^*, \omega^*)$;*

2. *$g_{aa} < 0$, $g_{ab} \leq 0$ and $g_{a\omega} > 0$.*

The first assumption rules out extreme misperceptions that result in the type of non-robustness studied in the previous application. The second assumption implies that the optimal effort level decreases in one's own effort and increases in the teammate's ability. Suppose the agent's initial model $\theta$ with parameter space $\Omega^\theta \subseteq [\underline{\omega}, \overline{\omega}]$ assigns probability 1 to $\hat{b} \in [\underline{b}, \overline{b}]$ and $\hat{b} \neq b^*$. That is, the agent is either dogmatically overconfident or underconfident about his own ability, but is ready to learn about his teammate's ability through his subjective model. Assumption 4 implies that beliefs are positively reinforcing when the agent is overconfident and negatively reinforcing when underconfident, as summarized by the following property.

**Property 1.** *Suppose $\hat{b} > b^*$ ($\hat{b} < b^*$). Then higher beliefs lead to higher optimal actions, i.e. $\max A^\theta(\delta_{\omega'}) \leq \min A^\theta(\delta_{\omega''})$ for all $\omega'' > \omega'$, and higher actions leads to higher (lower) beliefs, i.e. $\max \Omega^\theta(\delta_{a'}) \leq \min \Omega^\theta(\delta_{a''})$ ($\min \Omega^\theta(\delta_{a'}) \geq \max \Omega^\theta(\delta_{a''})$) for all $a'' > a'$.*

The first part that states higher beliefs induce higher actions immediately follows from the assumption that one's effort and his teammate's ability are complements. The intuition for the second part is as follows: as a result of overconfidence (underconfidence), he underestimates (overestimates) his teammate's ability; when a higher action is played, the return to the teammate's ability $\omega$ is higher because $g_{a\omega} > 0$, and thus the positive (negative) gap between the true state $\omega^*$ and the inferred teammate's ability $\hat{\omega}$ should be smaller such that expectations meet the reality, implying that the inferred state $\hat{\omega}$ is larger (smaller). Esponda et al. (2019) show that the agent's action frequency converges to a BN-E in this environment.

Consider any finite set of possible misperceptions $\mathcal{B} \subseteq [\underline{b}, \overline{b}]$. As in Application 1, to avoid trivial cases of non-robustness, I assume that for all $\hat{b} \in \mathcal{B}$ and all $a \in \mathcal{A}$, $\Omega^\theta$ contains the parameter value $\hat{\omega}$ if it satisfies $g(a, \hat{b}, \hat{\omega}) = g(a, b^*, \omega^*)$. Parameterizing the self-perception $b$, we have a $q$-family with a parameter space $\Omega^q = [\underline{b}, \overline{b}] \times [\underline{\omega}, \overline{\omega}]$. Then model $\theta$ belongs to the $q$-family if we augment its parameter space to $\overline{\Omega}^{\theta} = \{\hat{b}\} \times \Omega^{\theta}$. Proposition 2 characterizes the robustness properties for different self-perceptions. First, any level of overconfidence is globally robust but underconfidence is only guaranteed to be globally robust when it is small. Second, on unconnected intervals, underconfidence is either globally robust or not $q$-constrained locally robust in the sense that such a model does not persist against models with a self-perception closer to the true ability $b^*$.

**Proposition 2.** *There exists a strictly decreasing sequence $\{\beta_n\}_{1 \leq n \leq N}$ with $\beta_1 < b^*$, such that for any $\hat{b} \in \mathcal{B}$, the following statements hold:*

- *Models with overconfidence or mild underconfidence such that $\hat{b} > \beta_1$ are globally robust.*
- *Models with underconfidence such that $\hat{b} \in (\beta_{2k}, \beta_{2k-1})$ for some $k \in \mathbb{N}_+$ are not $q$-constrained locally robust. Given any such model $\theta$ and any $\tilde{b} \in \mathcal{B} \cap (\hat{b}, b^*]$, there exists a model $\theta^c \in \Theta^q$ associated with $\tilde{b}$ such that model $\theta$ does not persist against $\theta^c$.*
- *Models with underconfidence such that $\hat{b} \in (\beta_{2k+1}, \beta_{2k})$ for some $k \in \mathbb{N}_+$ are globally robust.*

The intuition behind Proposition 2 is as follows. If a Berk-Nash equilibrium $\sigma$ under model $\theta$ is pure, then it must be self-confirming because the agent can, by Assumption 4, perfectly justify his observations by forming an incorrect belief over his teammate's ability. By contrast, a mixed Berk-Nash equilibrium can never be self-confirming in this environment. By playing a mixed strategy, the agent necessarily finds his self-perception $\hat{b}$ inconsistent with the observations—no single value of teammate's ability can explain the outcome distributions at more than two effort levels when $\hat{b} \neq b^*$, and it can be shown that any $\tilde{b}$ that is slightly closer to the truth would explain the observations better. Therefore, Proposition 2 follows from Theorem 1 and Theorem 5, provided that we can show overconfidence ensures the existence of a pure and p-absorbing BN-E while underconfidence sometimes gives rise to mixed BN-Es only.

To illustrate the idea, suppose for simplicity that the agent chooses from three effort levels $a' < a^* < a''$. In particular, $a^*$ is the unique optimal effort level when the agent has a correct perception about the true DGP. Then by the continuity of $g$, when the agent's misperception is sufficiently small, $a^*$ is still a strict self-confirming equilibrium, supported by a slightly biased belief about his teammate's ability. Now consider an overconfident agent with $\hat{b} > b^*$ and suppose there arises a mixed BN-E in which the agent finds both $a'$ and $a^*$ optimal against a supporting belief $\delta_{\hat{\omega}}$. Then for any lower assessment of $\omega$, he strictly prefers the lower action $a'$, which, due to positive reinforcement, implies that he indeed finds a lower assessment of his teammate's ability more accurate. Therefore, the pure action $a'$ constitutes a uniformly strict SCE and the agent's subjective model is globally robust. Next, consider an underconfident agent and a potential mixed BN-E in which the agent finds both $a^*$ and $a''$ optimal against belief $\delta_{\hat{\omega}}$. Then for any higher assessment of $\omega$, he strictly prefers the higher action $a''$, which, due to negative reinforcement, induces his belief over $\omega$ to drift downwards; by contrast, for any lower assessment of $\omega$, he strictly prefers the lower action $a^*$, which induces his belief to drift upwards. Therefore, the mixed BN-E is the only BN-E and it follows that the agent's subjective model is not $q$-constrained locally robust. As $\hat{b}$ further decreases, the BN-E gradually changes from being mixed and not self-confirming to being pure and self-confirming, giving rise to the intervals as described in Proposition 2.

# 7 Concluding Remarks

In this paper, I develop and characterize three different robustness criteria of subjective models and use them to derive novel insights about which forms of model misspecification persist and which do not. Defined based on the chance of long-term persistence against competing models, the three robustness notions can be ranked in terms of how hard to satisfy their requirements and be used to compare different forms of model misspecification. If a model fails to be globally robust or unconstrained robust, it can still be constrained locally robust with respect to a particular parametric family. Further, a model may be constrained locally robust within one parametric family, but not so in a larger one.

Instead of assuming that the agent starts outright from a Berk-Nash equilibrium and compares how well models fit the data there, this framework incorporates model switching into full-fledged learning dynamics. The characterization highlights the importance of this consideration. Global robustness not only requires the existence of a self-confirming equilibrium but also needs it to be p-absorbing, which connects the notion of model robustness with the stability of equilibria under a single model. Further, global robustness requires that the agent must hold a relatively entrenched prior close to a self-confirming equilibrium belief and that she does not evaluate too many competing models at the same time. In a potential extension that involves persuasion by proposing competing models, these observations shed light on when and how such persuasion can be effective.

Within this general framework of model switching, there are many other interesting questions to pursue. For example, global robustness requires a positive chance of a model being adopted forever, but it may also be interesting to look at when a model is adopted infinitely often such that the underlying misspecification never vanishes. While the robustness approach delivers clean characterizations, one can restrict attention to a given small set of models and fully characterize the dynamics of model choices, i.e. how a decision maker oscillates between two or more competing models persistently.

# 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*.

Benoît, J.-P., J. Dubra, and D. A. Moore (2015): "Does the Better-than-average Effect Show that People are Overconfident?: Two Experiments," *Journal of the European Economic Association*, 13, 293–329.

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.

Brunnermeier, M. K. and J. A. Parker (2005): "Optimal expectations," *American Economic Review*, 95, 1092–1118.

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 (2021a): "Belief Convergence under Misspecified Learning: A Martingale Approach," .

——— (2021b): "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.

Köszegi, B. (2006): "Ego utility, overconfidence, and task choice," *Journal of the European Economic Association*, 4, 673–707.

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

Langer, E. J. and J. Roth (1975): "Heads I Win, Tails It's Chance: The Illusion of Control as a Function of the Sequence of Outcomes in a Purely Chance Task," *Journal of Personality and Social Psychology*, 32, 951.

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.

Oster, E., I. Shoulson, and E. Dorsey (2013): "Optimal expectations and limited medical testing: evidence from Huntington disease," *American Economic Review*, 103, 804–30.

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.

Svenson, O. (1981): "Are We All Less Risky and More Skillful Than Our Fellow Drivers?" *Acta psychologica*, 47, 143–148.

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.

# A Auxiliary Lemmas

**Lemma 4.** *Fix any $\theta \in \Theta$ and any correctly specified $\theta^c \in \Theta$. Consider a switcher whose initial model is $\theta$ and competing model is $\theta^c$, then as $t \to \infty$, $l_t^\theta / l_t^{\theta^c}$ a.s. converges to a non-negative random variable with finite expectation.*

*Proof.* Let $\kappa_t = l_t^\theta / l_t^{\theta^c}$, then $\kappa_0 = 1, \kappa_t \geq 0, \forall t$. I now construct the probability space in which $\kappa_t$ is a martingale. Given prior $\pi_0^{\theta^c}$, denote by $\mathbb{P}_S^{\theta^c}$ the probability measure over the set of histories $H$ as implied by model $\theta^c$. Formally, for any $\hat{H} \subseteq H$, we have $\mathbb{P}_S^{\theta^c}\left(\hat{H}\right) = \sum_{\omega \in \Omega^{\theta^c}} \pi_0^{\theta^c}(\omega) \mathbb{P}_S^{\theta^c,\omega}\left(\hat{H}\right)$, 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 $\theta^c$ and $\omega$. Take the conditional expectation of $\kappa_t$ with respect to $\mathbb{P}_S^{\theta^c}$, then we have

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

Hence, $\kappa_t$ is a martingale w.r.t. $\mathbb{P}_S^{\theta^c}$. Since $\kappa_t \geq 0, \forall t$, the Martingale Convergence Theorem implies that $\kappa_t$ converges to $\kappa$ almost surely w.r.t. $\mathbb{P}_S^{\theta^c}$, and $\mathbb{E}^{\mathbb{P}_S^{\theta^c}}\kappa \leq \mathbb{E}^{\mathbb{P}_S^{\theta^c}}\kappa_0 = 1$. Since $\theta^c$ 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 $\kappa_t$ also converges to $\kappa$ almost surely w.r.t. $\mathbb{P}_S^{\theta^c,\omega^*} \equiv \mathbb{P}_S$. Moreover, $\mathbb{E}\kappa < \infty$ because otherwise it contradicts $\mathbb{E}^{\mathbb{P}_S^{\theta^c}}\kappa \leq 1$. $\square$

**Lemma 5.** *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*

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

*Then for any fixed $t_0 \geq 1$,*

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

*Proof.* $\xi_t (h)$ is a martingale difference process since $E(\xi_t (h) | h_t) = 0$. So for any $t_0$, $\xi_{t_0}^t (h) := \sum_{\tau=t_0}^{t} (t - \tau + 1)^{-1} \xi_\tau (h)$ 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

$$\begin{aligned}
\mathbb{E}\left(\left(\xi_{t_0}^t\right)^2\right) =& \mathbb{E}\left[\left(\sum_{\tau=t_0}^{t} (t - \tau + 1)^{-1} \xi_\tau (h)\right)^2\right] \\
\leq& \sum_{\tau=t_0}^{t} (t - \tau + 1)^{-2} \mathbb{E}\left[(\xi_\tau (h))^2\right] \\
\leq& \sum_{\tau=t_0}^{t} (t - \tau + 1)^{-2} \mathbb{E}\left[\left(\ln \frac{q^\theta (y_t | a_t, \omega)}{q^{\theta'} (y_t | a_t, \omega')}\right)^2\right] \\
\leq& 2 \sum_{\tau=t_0}^{t} (t - \tau + 1)^{-2} \mathbb{E}\left[\left(\ln \frac{q^* (y_t | a_t)}{q^\theta (y_t | a_t, \omega)}\right)^2 + \left(\ln \frac{q^* (y_t | a_t)}{q^{\theta'} (y_t | a_t, \omega')}\right)^2\right] \\
\leq& 4 \sum_{\tau=t_0}^{t} (t - \tau + 1)^{-2} \max_a \mathbb{E}\left[(g_a (y))^2\right] < \infty,
\end{aligned}$$

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 $\mathbb{E}\left(\left(\xi_{t_0}^\infty\right)^2\right) < \infty$. 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., it follows from the Kronecker Lemma that

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

□

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 6.** *Suppose the action frequency of a dogmatic modeler with $\Theta^{\dagger} = \{\theta\}$ converges to $\sigma$, 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 $\sigma$, then her belief $\pi_t^{\theta}$ also a.s. converges to some $\pi^{\theta}$ 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 $\Omega^{\theta}$ is finite, for any given $\sigma$, there exists $\epsilon > 0$ such that

$$\sum_{\mathcal{A}} \sigma(a)\left[D_{KL}\left(q^*(\cdot|a) \parallel q^{\theta}(\cdot|a,\omega)\right) - D_{KL}\left(q^*(\cdot|a) \parallel q^{\theta}(\cdot|a,\omega')\right)\right] < -\epsilon \tag{16}$$

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. (16) as $\sigma_t$ converges to $\sigma$. 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 5, 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. $\square$

**Lemma 7.** *Fix any $\theta \in \Theta$, the optimal action correspondence $A^\theta : \Delta\Omega^\theta \rightrightarrows \mathcal{A}$ is upper hemicontinuous in both the belief $\pi$ and the discount factor $\delta$.*

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

**Lemma 8.** *Fix any $\theta \in \Theta$, the set of all Berk-Nash equilibria under $\theta$ is compact.*

*Proof.* Denote the set of all Berk-Nash equilibria under model $\theta$ as $BN^\theta \subseteq \Delta\mathcal{A}$. Since $\Delta\mathcal{A}$ is bounded, we only need to show that $BN^\theta$ is closed. Suppose $\sigma$ is the limit of some sequence $(\sigma_n)_n$ of Berk-Nash equilibria, but $\sigma$ 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 $\Omega^\theta(\sigma_n) \subseteq \Omega^\theta(\sigma)$ 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 $\operatorname{supp}(\sigma) \subseteq \operatorname{supp}(\sigma_n)$ for large enough $n$, which implies that $\sigma_n \notin \Delta A_m^\theta(\pi)$ for large $n$. This is a contradiction. $\square$

# B Proofs of Results

## B.1 Proof of Theorem 1

I first prove Lemmas 1, 2 and 3.

*Proof of Lemma 1.* It immediately follows from Lemma 4 that $l_t^{\theta^c}/l_t^\theta$ a.s. converges to $\iota \leq \alpha$ on paths where $m_t$ converges to $\theta$. I now show that $\pi_t^\theta$ 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 a model that consists of a single parameter $\omega$ and the competing model $\theta^c$. By Lemma 4, $l_t^{\theta,\omega}/l_t^{\theta^c}$ a.s. converges to a random variable with finite expectation. Consider the paths on which $m_t$ converges to $\theta$. 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 to a random variable with finite expectation as well. Since this is true for all $\omega \in \Omega^{\theta}$, $\pi_t^{\theta}$ a.s. converges to some limit $\pi_{\infty}^{\theta}$ 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 4, converges almost surely. □

*Proof of Lemma 2.* Consider any $\hat{\omega}$ such that with positive probability, $m_t$ converges to $\theta$ and $\hat{\omega} \in \operatorname{supp}(\pi_{\infty}^{\theta})$. Let $A^-(\hat{\omega}) \equiv \{a \in \mathcal{A} : q^{\theta}(\cdot|a,\hat{\omega}) \neq q^*(\cdot|a)\}$. 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 $\theta$. 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 $\theta^c$ 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, by Lemma 5, a.s. increases to infinity as $t \to \infty$, contradicting the assumption that $m_t$ converges to $\theta$. Therefore, on the paths where $m_t \to \theta$, almost surely, there exists $T$ such that $a_t \in \mathcal{A}\setminus \cup_{\omega'\in\operatorname{supp}(\pi_{\infty}^{\theta})} A^-(\hat{\omega}), \forall t > T$.

Since $q^{\theta}(\cdot|a,\omega') \equiv q^*(\cdot|a)$ for all $\omega' \in \operatorname{supp}(\pi_{\infty}^{\theta})$ and all $a \in \mathcal{A}\setminus \cup_{\omega'\in\operatorname{supp}(\pi_{\infty}^{\theta})} A^-(\omega')$, the actions that are played in the limit have vanishing experimentation value and thus are myopically optimal. Therefore, fixing a particular value of $\pi_{\infty}^{\theta}$ that is a limit belief for a positive measure of histories where $m_t \to \theta$, there exists a set of actions $\hat{A} \subseteq A_m^{\theta}(\pi_{\infty}^{\theta})$ on those histories, the agent only plays actions from this set in the limit. Since $m_t$ eventually converges to $\theta$, it must be true that with positive probability, a dogmatic modeler who inherits the switcher's prior and policy from the period when the last switch happens also only plays actions from $\hat{A}$ in the limit. Therefore, any strategy $\sigma$ with $\operatorname{supp}(\sigma) = \hat{A}$ is a p-absorbing self-confirming equilibrium under $\theta$. □

*Proof of Lemma 3.* Suppose there exists a p-absorbing SCE $\sigma$ under $\theta$. Consider the learning process of a dogmatic modeler with model $\theta$. There exists a full-support prior $\pi_0^\theta \in \Delta\Omega^\theta$ and an optimal policy $a^\theta$ 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 a SCE $\sigma'$ 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, $\pi_t^\theta$ a.s. converges to a limit $\pi_\infty^\theta$ on $\tilde{H}$, with $\text{supp}\left(\pi_\infty^\theta\right) \subseteq \Omega^\theta(\sigma) = \{\omega \in \Omega^\theta : q^*\left(\cdot|a\right) = q^\theta\left(\cdot|a,\omega\right), \forall a \in \text{supp}\,(\sigma)\}$.

This implies the existence of an integer $T > 0$ such that, with positive probability, we have (1) $a_t \in \text{supp}(\sigma), \forall t \geq T$, (2) $\pi_t^\theta$ converges to a limit $\pi_\infty^\theta$ with $\text{supp}(\pi_\infty^\theta) \subseteq \Omega^\theta(\sigma)$. Let $\epsilon$ be any positive constant. We can find a posterior a new prior $\tilde{\pi}_0^\theta \in B_\epsilon(\Delta\Omega^\theta(\sigma))$ under which, on a positive measure of histories, a dogmatic modeler behaves such that (1′) $a_t \in \text{supp}(\sigma), \forall t \geq 0$, and (2′) the posterior $\hat{\pi}_t^\theta$ a.s. converges to $\pi_\infty^\theta$ and never leaves $B_\epsilon(\Delta\Omega^\theta(\sigma)), \forall t \geq 0$.

Denote as $E$ the event described by (1′) and (2′). I now show for any constant $\gamma \in (0,1)$, there exists a full-support prior $\hat{\pi}_0^\theta$ under which $E$ occurs with probability higher than this constant. Suppose for contradiction that this is not true. Denote the probability of $E$ under any full-support prior by $\gamma(\pi^\theta)$ and let $\overline{\gamma} := \sup_{\pi_0^\theta \in \Delta\Omega^\theta} \gamma(\pi^\theta)$. By assumption, $\overline{\gamma} < 1$. By definition, for any small $\psi \in (0,1)$, there exists some prior $\pi_0^{\theta,\psi}$ such that $\gamma(\pi_0^{\theta,\psi}) > \overline{\gamma} - \psi$. Notice that with probability $1 - \pi_0^{\theta,\psi}$, the dogmatic modeler eventually either arrives at some posterior $\pi_t^{\theta,\psi}$ such that it either leads her to play an action outside $\text{supp}(\sigma)$ or is itself outside $B_\epsilon(\Delta\Omega^\theta(\sigma))$. Hence, there exists an integer $T > 0$ such that

$$Pr\left(\gamma(\pi_T^{\theta,\psi}) = 0\right) > \gamma(\pi_0^{\theta,\psi}) - \psi > \overline{\gamma} - 2\psi.$$

Now, consider the maximum probability that $E$ is achieved if the agent starts with a prior that is equal to one of the possible posteriors $\pi_T^{\theta,\psi}$. Since

$$\gamma(\pi_0^{\theta,\psi}) = \mathbb{E}_{h_T \in H_T} \gamma(\pi_T^{\theta,\psi}),$$

we have

$$\begin{aligned}\max_{h_T \in H_T} \gamma(\pi_T^{\theta,\psi}) &> \frac{\gamma(\pi_0^{\theta,\psi})}{1 - Pr\left(\gamma(\pi_T^{\theta,\psi}) = 0\right)} \\ &> \frac{\overline{\gamma} - \psi}{1 - \overline{\gamma} + 2\psi}.\end{aligned}$$

But notice that when $\psi$ is sufficiently small, the term $\frac{\overline{\gamma}-\psi}{1-\overline{\gamma}+2\psi} > \overline{\gamma}$, contradicting the assumption that $\overline{\gamma}$ is the supremum of $\gamma(\pi^\theta)$ over all full-support beliefs.

$\square$

Now take any competing model $\theta'$ with any parameter space $\Omega^{\theta'}$ and predictions $q^{\theta'}$, and any full-support prior $\pi_0^{\theta'}$ and optimal policy $a^{\theta'}$. Define

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

which can be shown to be a martingale with respect to both $\mathbb{P}_B$ and $\mathbb{P}_S$ by a similar argument as in Lemma 4. By the Ville's maximal inequality for supermartingales, the probability that $S_n$ is bounded above by a positive constant is bounded. In particular, for any $\eta \in (1, \alpha)$,

$$\mathbb{P}_B(S_t \leq \eta, \forall t \geq 0) \geq 1 - \frac{\mathbb{E}^{\mathbb{P}_B} S_0}{\eta} = 1 - \frac{1}{\eta}.$$

Note that this inequality holds regardless of the prior and policy assigned to models $\theta$ and $\theta'$ and the structure of model $\theta'$.

By Lemma 3, we know that for any $\eta \in (1, \alpha)$ and $\epsilon > 0$, there exist a prior $\pi_0^\theta$ and an optimal $a^\theta$ such that the probability of event $E$ (as defined in the proof of Lemma 3) exceeds $1/\eta$. Combining this with the previous observation, we know that the probability that (1′) and (2′) hold and $S_t$ never exceeds $\eta$ at the same time must be positive. Denote the infinite histories that these events occur by $\hat{H}$. Consider the likelihood ratio $l_t^{\theta'}/l_t^\theta$ computed from any history in $\hat{H}$ under prior $\pi_0^{\theta,\psi}$. Then when $\epsilon$ is small enough,

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

where the first inequality follows from the fact that $\pi_0^\theta$ is full-support, the second inequality follows from the definition of $\hat{H}$. Thus, on histories $\hat{H}$, the switcher never makes any switch to the competing model model $\theta' \in \Theta^c$, i.e. $m_t = \theta, \forall t \geq 0$. Therefore, if we endow the switcher with the same prior $\pi_0^\theta$, 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 $\sigma$ is p-absorbing. By definition, there exists a belief $\pi \in \Delta\Omega^\theta$ with $\operatorname{supp}(\pi) \subseteq G_\sigma^\theta$ (the set $G_\sigma^\theta$ is defined in the proof of Theorem 1). Since $\sigma$ is uniformly quasi-strict, $\text{supp}(\sigma)$ contains all myopically optimal actions against each degenerate belief $\delta_\omega$ concentrated on $\omega \in \text{supp}(\pi)$. In addition, $\text{supp}(\sigma)$ must be optimal against $\delta_\omega$ for an agent who maximizes discounted utility, because the dynamic programming problem described by (5) reduces to a static maximization problem when the belief is degenerate. This implies that $\text{supp}(\sigma)$ is also optimal against $\pi$. Further, since $A^\theta$ is upper hemicontinuous (by Lemma 7), there exists $\gamma > 0$ small enough such that $\text{supp}(\sigma) = A^\theta(\tilde{\pi})$ for all $\tilde{\pi} \in B_\gamma(\pi)$.

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

$$\begin{aligned}
\mathbb{E}\left[\frac{\pi^\theta_t(\omega)}{\pi^\theta_t(G^\theta_\sigma)}|h_t\right] &= \mathbb{E}\left[\frac{\pi^\theta_0(\omega)\prod_{\tau=0}^{t-1} q^\theta(y_\tau|a_\tau,\omega)}{\sum_{\omega'\in G^\theta_\sigma}\pi^\theta_0(\omega')\prod_{\tau=0}^{t-1} q^\theta(y_\tau|a_\tau,\omega')}|h_t\right] \\
&= \mathbb{E}\left[\frac{\pi^\theta_0(\omega)}{\pi^\theta_0(G^\theta_\sigma)}\frac{\prod_{\tau=0}^{t-1} q^\theta(y_\tau|a_\tau,\omega)}{\prod_{\tau=0}^{t-1} q^*(y_\tau|a_\tau)}|h_t\right] \\
&= \frac{\pi^\theta_0(\omega)\prod_{\tau=0}^{t-2} q^\theta(y_\tau|a_\tau,\omega)}{\pi^\theta_0(G^\theta_\sigma)\prod_{\tau=0}^{t-2} q^*(y_\tau|a_\tau)} = \frac{\pi^\theta_{t-1}(\omega)}{\pi^\theta_{t-1}(G^\theta_\sigma)}
\end{aligned}$$

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

$$\mathbb{P}_B\left(\frac{\pi^\theta_t(\Omega^\theta\backslash G^\theta_\sigma)}{\pi^\theta_t(G^\theta_\sigma)} \geq \frac{\pi^\theta_0(\Omega^\theta\backslash G^\theta_\sigma)}{\pi^\theta_0(G^\theta_\sigma)} + \epsilon \text{ for some } t\right) < 1.$$

Since $\pi^\theta_t(G^\theta_\sigma) = 1 - \pi^\theta_t(\Omega^\theta\backslash G^\theta_\sigma)$, the above inequality implies that for all $\epsilon > 0$,

$$\mathbb{P}_B\left(\pi^\theta_t(\Omega^\theta\backslash G^\theta_\sigma) \geq \pi^\theta_0(\Omega^\theta\backslash G^\theta_\sigma) + \epsilon \text{ for some } t\right) < 1.$$

Pick some $\gamma' \in (0,\gamma)$ and $\pi^\theta_0 \in B_{\gamma'}(\pi)$, then $\pi^\theta_0(G^\theta_\sigma) > 1-\gamma'$. Notice that the belief ratio $\frac{\pi^\theta_t(\omega)}{\pi^\theta_t(\omega')}$ remain unchanged throughout all periods provided that $\omega,\omega' \in G^\theta_\sigma$. Hence, if $\pi^\theta_t \notin B_\gamma(\pi)$ for some $t \geq 0$, then there exists $t$ such that $\pi^\theta_t(\Omega^\theta\backslash G^\theta_\sigma) \geq \pi^\theta_0(\Omega^\theta\backslash G^\theta_\sigma) + \gamma - \gamma'$. Therefore,

$$\begin{aligned}
&\mathbb{P}_B\left(\pi^\theta_t \notin B_\gamma(\pi) \text{ for some } t \geq 0\right) \\
\leq &\mathbb{P}_B\left(\pi^\theta_t(\Omega^\theta\backslash G^\theta_\sigma) \geq \pi^\theta_0(\Omega^\theta\backslash G^\theta_\sigma) + \gamma - \gamma' \text{ for some } t\right) < 1.
\end{aligned}$$

This implies that $\mathbb{P}_B\left(\pi^\theta_t \in B_\gamma(\pi), \forall t \geq 0\right) > 0$. Notice that $\pi^\theta_t \in B_\gamma(\pi), \forall t \geq 0$ in turn implies that $a_t \in \text{supp}(\sigma), \forall t \geq 0$. Therefore, $\sigma$ is p-absorbing.

## B.3 Proof of Theorem 3

I now show that if $\theta$ 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 $\theta$ as $\Omega^{\theta} = \left\{\omega^1, ..., \omega^N\right\}$. Consider a competing model $\theta^c$ constructed as below:

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

By construction, $\theta^c \in N_{\epsilon}(\theta)$. Hence, there exists $\epsilon > 0$ such that $\theta$ persists against $\theta^c$ under some full-support priors $\pi_0^{\theta}$, $\pi_0^{\theta^c}$ and some policies $a^{\theta}, a^{\theta^c}$. 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

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

Notice that the last term, denoted by $l_t^*/l_t^{\theta}$, is the likelihood ratio of the true DGP and $\theta$. 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, $\theta$ 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 $\theta$ is globally robust.

## B.4 Proof of Theorem 4

Suppose $\sigma$ is a pure p-absorbing BN-E with $\theta$ being locally KL-minimizing and locally identified at $\sigma$ w.r.t. $\Omega^q$, and $\sigma$ assigns probability 1 to $a^* \in \mathcal{A}$. Then there exists a full-support prior $\pi_0^{\theta}$ and a policy $a^{\theta}$ such that a dogmatic modeler eventually only plays $a^*$ with positive probability. It follows from Lemma 6 that $\pi_t^{\theta}\left(\Omega^{\theta}(\sigma)\right) \xrightarrow{a.s.} 1$. Fix any $\gamma \in (0,1)$. By a similar argument as in Lemma 3, we know that for any $\eta \in (1, \alpha)$, there exists a full-support prior $\pi_0^{\theta}$ and a policy under which with probability larger than $1/\eta$, we have $a_t \in \text{supp}(\sigma)$ and $\pi_t^{\theta}\left(\Omega^{\theta}(\sigma)\right) > \gamma, \forall t \geq 0$.

Since $\Omega^{\theta}$ is finite, there exists $\epsilon_1$ such that whenever $\epsilon < \epsilon_1$, the $\epsilon$-neighborhoods of any two $\omega$ and $\tilde{\omega} \in \Omega^{\theta}$ do not coincide. It follows from Lemma 2 in Frick et al. (2021a) that the local dominant condition implies that all different KL minimizers in $\Omega^\theta(\sigma)$ correspond to the same distribution under $\sigma$. Moreover, there exists $\epsilon_2 > 0$ s.t. whenever $\epsilon < \epsilon_2$, we have

$$\mathbb{E}\left(\frac{q(\cdot|a,\omega')}{q(\cdot|a,\omega)}\right)^d < 1$$

for all $\omega \in \Omega^\theta(\sigma)$, all $\omega' \in \Omega^q \cap B_\epsilon(\Omega^\theta(\sigma)) \setminus \Omega^\theta(\sigma)$, and all $a \in \text{supp}(\sigma)$.

Let $\epsilon < \min\{\epsilon_1, \epsilon_2\}$. Pick any competing model $\theta^c \in N^q_\epsilon(\theta)$ with belief $\pi_0^{\theta^c} \in D^q_\epsilon(\pi_0^\theta)$. Then there exists a correspondence $\iota : \Omega^\theta \rightrightarrows \Omega^{\theta^c}$ s.t. $\iota(\omega) = \Omega^{\theta^c} \cap B_\epsilon(\omega) \setminus \Omega^\theta(\sigma)$ and $\cup_{\Omega^\theta}\iota(\omega) \equiv \Omega^{\theta^c} \setminus \Omega^\theta(\sigma)$. Pick any $\omega^* \in \Omega^\theta(\sigma)$. Since $q$ is uniformly continuous in $\omega$, there exists $\epsilon_3 \le \epsilon_2$ such that if $\epsilon < \epsilon_3$, then at all actions $a \in \text{supp}(\sigma)$, we have

$$\mathbb{E}\left(\frac{\sum_{\omega' \in \iota(\omega)} f(\omega')q(\cdot|a,\omega')}{q(\cdot|a,\omega^*)}\right)^d < 1$$

for all $\omega \in \Omega^\theta \setminus \Omega^\theta(\sigma)$ and all probability distributions $f$ over $\iota(\omega)$, and

$$\mathbb{E}\left(\frac{\sum_{\omega' \in \cup_{\Omega^\theta(\sigma)}\iota(\omega)} f(\omega')q(\cdot|a,\omega')}{q(\cdot|a,\omega^*)}\right)^d < 1$$

for all probability distributions $f$ over $\cup_{\Omega^\theta(\sigma)}\iota(\omega)$.

Define $l_t^{q,\omega} := \prod_{\tau=0}^t q(y_\tau|a_\tau,\omega)$. Let $\xi_t$ be an indicator function such that $\xi_t = 1$ if $a_t \in \text{supp}(\sigma)$ and $\xi_t = 0$ otherwise. Note that $\xi_t$ is measurable w.r.t. the measure over $H_{t-1}$. Define a new distribution of outcomes $\hat{q}$, s.t.

$$\hat{q}(\cdot|a_t,\omega^*) := \xi_t q(\cdot|a_t,\omega^*) + (1-\xi_t)q^*(\cdot|a_t).$$

Let $\hat{l}^{q,\omega^*} := \prod_{\tau=0}^t \hat{q}(y_\tau|a_\tau,\omega^*)$. For any $\omega \in \Omega^\theta \setminus \Omega^\theta(\sigma)$ and any distribution $f$ over $\iota(\omega)$,

$$\begin{aligned}
&\mathbb{E}^{\mathbb{P}_B}\left(\left(\frac{\sum_{\omega'\in\iota(\omega)} f(\omega')l_t^{q,\omega'}}{\hat{l}_t^{q,\omega^*}}\right)^d |h_{t-1}\right) \\
&= \left(\frac{\sum_{\omega'\in\iota(\omega)} f(\omega')l_{t-1}^{q,\omega'}}{\hat{l}_{t-1}^{q,\omega^*}}\right)^d \mathbb{E}^{\mathbb{P}_B}\left(\left(\frac{\sum_{\omega'\in\iota(\omega)} \hat{f}(\omega')q(y_t|a_t,\omega')}{\hat{q}(y_t|a_t,\omega^*)}\right)^d |h_{t-1}\right) \\
&< \left(\frac{\sum_{\omega'\in\iota(\omega)} f(\omega')l_{t-1}^{q,\omega'}}{\hat{l}_{t-1}^{q,\omega^*}}\right)^d,
\end{aligned}$$

where $\hat{f}(\omega') = f(\omega')l_{t-1}^{q,\omega'}/\left(\sum_{\omega''\in\iota(\omega)} f(\omega'')l_{t-1}^{q,\omega''}\right)$. The inequality follows from our previous observation. Hence, by Ville's maximal inequality for supermartingales, we know that for any $\kappa > 1$,

$$\mathbb{P}_B\left(\frac{\sum_{\omega'\in\iota(\omega)} f(\omega')l_t^{q,\omega'}}{\hat{l}_t^{q,\omega^*}} \leq \kappa, \forall t\right) \geq 1 - \frac{1}{\kappa^{1/d}}.$$

Similarly, for all probability distributions $f$ over $\cup_{\Omega^\theta(\sigma)}\iota(\omega)$ and any $\eta \in (0, \alpha)$, we have that

$$\mathbb{P}_B\left(\frac{\sum_{\omega'\in\cup_{\Omega^\theta(\sigma)}} f(\omega')l_t^{q,\omega'}}{\hat{l}_t^{q,\omega^*}} \leq \eta, \forall t\right) \geq 1 - \frac{1}{\eta^{1/d}}.$$

When $\kappa$ is sufficiently large, we have

$$\begin{aligned}
&\mathbb{P}_B\left(\frac{\sum_{\omega^c\in\Omega^{\theta^c}} \pi^{\theta^c}(\omega^c)l_t^{q,\omega^c}}{\hat{l}_t^{q,\omega^*}} \leq \pi^{\theta^c}(\Omega^\theta(\sigma))\eta + \pi^{\theta^c}(\Omega^\theta \setminus \Omega^\theta(\sigma))\kappa, \forall t \geq 0\right) \\
\geq& 1 - \frac{1}{\eta^{1/d}} + M \cdot \left(1 - \frac{1}{\kappa^{1/d}}\right) - M = 1 - \frac{1}{\eta^{1/d}} - \frac{M}{\kappa^{1/d}} > 0.
\end{aligned} \tag{17}$$

Note that when $\pi^{\theta^c} \in D_\epsilon^q(\pi^\theta)$ and $\pi^\theta(\Omega^\theta(\sigma)) > 1 - \epsilon$, we have

$$\pi^{\theta^c}(\Omega^\theta(\sigma))\eta + \pi^{\theta^c}(\Omega^\theta \setminus \Omega^\theta(\sigma))\kappa \leq \eta + 2\epsilon\kappa.$$

Also, there exists a prior $\pi_0^\theta$ under which with probability larger than $1 - \frac{1}{\eta^{1/d}} - \frac{M}{\kappa^{1/d}}$, the dogmatic modeler's behavior satisfies that $a_t \in \text{supp}\,(\sigma)$ and $\pi_t^\theta\left(\Omega^\theta\left(\sigma\right)\right) > 1 - \epsilon, \forall t \geq 0$. In summary, when $\epsilon < \min\{\epsilon_1, \epsilon_2, \epsilon_3\}$, there exists a prior $\pi_0^\theta$ such that the following event happens with positive probability:

$$\begin{gathered}
a_t \in \text{supp}\,(\sigma) \text{ and } \pi_t^\theta\left(\Omega^\theta\left(\sigma\right)\right) > 1 - \epsilon, \forall t \geq 0 \\
\frac{\sum_{\omega^c\in\Omega^{\theta^c}} \pi^{\theta^c}(\omega^c)l_t^{q,\omega^c}}{l_t^{q,\omega^*}} \leq \eta + 2\epsilon\kappa, \forall t \geq 0
\end{gathered}$$

When $\epsilon$ is small enough, conditional on the above event, we have

$$\frac{l_t^{\theta^c}}{l_t^\theta} < \frac{\sum_{\omega^c\in\Omega^{\theta^c}} \pi^{\theta^c}(\omega^c)l_t^{q,\omega^c}}{\pi_0^\theta(\Omega^\theta(\sigma))l_t^{q,\omega^*}} < \frac{\eta + 2\epsilon\kappa}{1 - \epsilon} < \alpha.$$

Hence, conditional on this event, the switcher never switches to the competing model $\theta^c$.

## B.5 Proof of Theorem 5

Suppose that there exist no Berk-Nash equilibrium $\sigma$ under $\theta$ with $\theta$ being locally KL-minimizing at $\sigma$ w.r.t. $\Omega^q$. Suppose for the sake of contradiction that $\theta$ is $q$-constrained locally robust within a neighborhood of $\epsilon$.

Define a function $K^q : \Delta\mathcal{A} \times \Omega^p \to \mathbb{R}$, where

$$K^q(\sigma, \omega) \coloneqq \sum_{\mathcal{A}} \sigma(a) D_{KL}(q^*(\cdot|a) \parallel q(\cdot|a, \omega)). \tag{18}$$

That is, $K^q(\sigma, \omega)$ represents the $\sigma$-weighted KL divergence between the outcome distribution predicted by $\omega$ and the true DGP.

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

$$\min_{\omega \in \Omega^\theta} K^q(\sigma, \omega) > K^q(\sigma, \omega'). \tag{19}$$

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

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

We know from Lemma 8 that the set of Berk-Nash equilibria under $\theta$ is compact. Therefore, by the Heine-Borel theorem, there must exist finite number of parameters, collected by a set $R_\epsilon$, such that for any Berk-Nash equilibrium $\sigma$, we can find some parameter from the set $R_\epsilon$ such that the above inequality (19) holds.

Consider a competing model $\theta^c$ 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 $\epsilon$ evenly to $R_\epsilon$. Formally, let

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

Consider all possible histories in which the switcher eventually adopts $\theta$. Then the switcher's action frequency a.s. converges to a Berk-Nash equilibrium by assumption. Consider the paths where this limit equilibrium is $\sigma$. 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

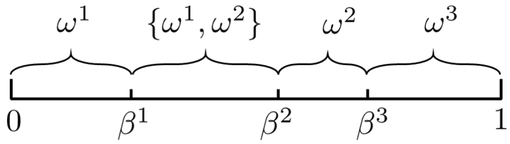


Figure 1: Example of a binary-action setting: Each point in the interval represents a mixed action's weight assigned to $a^2$; the parameter(s) placed above a segment of the interval are the minimizer(s) of $K^q(\sigma, \omega)$ in $\Omega^\theta$ for all $\sigma$ in this segment.

$\omega'' \in R_\epsilon$ such that $K^q(\sigma_t, \omega'') - K^q(\sigma_t, \omega) < -\eta, \forall \omega \in \Omega^\theta$. It then follows that

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

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

## B.6 A Necessary Condition for Constrained Local Robustness

The following theorem states a necessary condition for constrained local robustness in a special environment with binary actions and a myopic agent.

**Theorem 7.** *Suppose $|\mathcal{A}| = 2$ and $\delta = 0$. Then a model $\theta \in \Theta^q$ is $q$-constrained locally robust only if only if it admits a BN-E $\sigma$ at which $\theta$ is locally KL-minimizing w.r.t. $\Omega^q$.*

The critical step in proving Theorem 7 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 5. 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 $\Omega^\theta$ has four elements, each of which is a KL minimizer in $\Omega^\theta$ 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, $\omega^1$ uniquely minimizes the KL divergence when evaluated at a mixed action when $\beta \in [0, \beta^1]$, while both $\omega^1$ and $\omega^2$ 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 $\omega^1$ and $\omega^2$ are KL minimizers infinitely often but not $\omega^3$. Since the action space is binary, if $\sigma_t$ enters two non-connected regions on the unit interval infinitely often, it must also cross the region in between infinitely often.[18] This implies that $\sigma_t$ must enter $[\beta^1, \beta^2]$ infinitely often. To generate this pattern, it must be that $a^2 \in A_m^\theta(\delta_{\omega^1})$ and $a^1 \in A_m^\theta(\delta_{\omega^2})$, because otherwise only one action will be played in the limit.[19] Thus, there exists a mixed belief over $\omega^1$ and $\omega^2$ that makes the myopic agent indifferent between the actions. Since both $\omega^1$ and $\omega^2$ 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.

*Proof.* We only need to show that given any $\epsilon > 0$, almost surely, a dogmatic modeler's action frequency $\sigma_t$ enters the $\epsilon$-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 5, it can be shown that $\theta$ is not $q$-constrained locally robust if there is no Berk-Nash equilibrium $\sigma$ such that $\theta$ is locally KL-minimizing at $\sigma$.

For convenience, let $\mathcal{A} = \{a^1, a^2\}$. First, consider the paths where $\sigma_t$ converges to some limit $\sigma$. denoted by $H^1$. Then Lemma 6 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 $\sigma$. Hence, for each action $a$ in the support of $\sigma$, there exists some belief $\pi_a \in \Delta\Omega^\theta(\sigma)$ such that $a \in A_m^\theta(\pi_a)$. If $\operatorname{supp}(\sigma)$ is a singleton, then this immediately implies that $\sigma$ is a Berk-Nash equilibrium. If instead $\operatorname{supp}(\sigma) = \{a^1, a^2\}$, then by the hemi-continuity of $A_m^\theta$, 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 $\sigma$ is a Berk-Nash equilibrium. Therefore, her action frequency $\sigma_t$ enters the $\epsilon$-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 $\Omega_\infty^\theta$ be the set of all parameters in $\Omega^\theta$ 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 $\sigma_\omega$ that satisfies $\omega \in \Omega^\theta(\sigma_\omega)$ is a Berk-Nash equilibrium.

[18] 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.

[19] $\delta_{\omega^1}$ and $\delta_{\omega^2}$ denote the degenerate beliefs at $\omega^1$ and $\omega^2$, respectively.

By construction, this means $\sigma_t$ 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 $\sigma_t$ oscillates, $\Omega^{\theta}_{\infty}$ 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 $\omega$ and $\omega' \in \Omega^{\theta}_{\infty}$, say they are *related* if there exists some mixed action $\sigma$ 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 $\Omega^{\theta}_{\infty}$ must be related to some other parameter in $\Omega^{\theta}_{\infty}$. 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 $\gamma$ such that $B_{\gamma}(C_{\omega^*}) \cap \left(\cup_{\omega \in \Omega^{\theta}_{\infty} \setminus \{\omega^*\}} C_{\omega}\right) = \emptyset$. However, since $\omega^*$ 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 $\sigma_t$ infinitely often yet not included by $\Omega^{\theta}_{\infty}$, contradicting the definition of $\Omega^{\theta}_{\infty}$. By the same logic, there cannot be two cliques in $\Omega^{\theta}_{\infty}$ 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 $\Omega^{\theta}_{\infty}$ 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 $C_{\omega}$ is compact and convex. Since $\sigma_t$ enters both $C_{\omega}$ and $C_{\omega'}$ infinitely many times, it must be that $\sigma_t$ enters the $\epsilon$-neighborhood of $C_{\omega} \cap C_{\omega'}$ infinitely often for any $\epsilon > 0$. The proof is now complete. □

## B.7 Proof of Theorem 6

To show that Theorem 1 continues to hold when $\alpha > K$, it suffices to show that a model $\theta$ is globally robust if $\theta$ admits a p-absorbing SCE. Without loss of generality, take any $\Theta^c = \theta^1, ..., \theta^K \subseteq \Theta$ and define for each $k \in \{1, ..., K\}$,

$$S^k_t = \frac{\sum_{\omega' \in \Omega^{\theta^k}} \pi^{\theta^k}_0(\omega') \prod_{\tau=0}^{t} q^{\theta^k}(y_\tau | a_\tau, \omega')}{\prod_{\tau=0}^{t} q^*(y_\tau | a_\tau)}.$$

Then for any $\eta \in (1, \alpha)$, we have

$$\mathbb{P}_B(S^k_t \leq \eta, \forall t \geq 0) \geq 1 - \frac{\mathbb{E}^{\mathbb{P}_B} S^k_0}{\eta} = 1 - \frac{1}{\eta}.$$

Hence, when $\eta$ is sufficiently close to $\alpha$,

$$
\begin{aligned}
&\mathbb{P}_B(S_t^k \leq \eta, \forall t \geq 0, \forall k \in \{1, ..., K\}) \\
\geq &1 - \sum_{k=1}^{K} P_B(S_t^k > \eta \text{ for some } t \geq 0) \\
\geq &1 - \frac{K}{\eta} > 0.
\end{aligned}
$$

The rest of the argument is identical to the proof in Section B.1. It follows that Theorem 3 also continues to hold when $\alpha > K$.

For $q$-constrained local robustness, note that Inequality (17) implies

$$
\begin{aligned}
&\mathbb{P}_B\left(\frac{\sum_{\omega^k \in \Omega^{\theta^k}} \pi^{\theta^k}(\omega^k) l_t^{q,\omega^k}}{\hat{l}_t^{q,\omega^*}} \leq \pi^{\theta^k}(\Omega^\theta(\sigma))\eta + \pi^{\theta^k}(\Omega^\theta \setminus \Omega^\theta(\sigma))\kappa, \forall t \geq 0, \forall k \in \{1, ..., K\}\right) \\
\geq &1 - \sum_{k=1}^{K} P_B\left(\frac{\sum_{\omega^k \in \Omega^{\theta^k}} \pi^{\theta^k}(\omega^k) l_t^{q,\omega^k}}{\hat{l}_t^{q,\omega^*}} \leq \pi^{\theta^k}(\Omega^\theta(\sigma))\eta + \pi^{\theta^k}(\Omega^\theta \setminus \Omega^\theta(\sigma))\kappa \text{ for some } t \geq 0\right) \\
\geq &1 - K\left(\frac{1}{\eta^{1/d}} + \frac{M}{\kappa^{1/d}}\right).
\end{aligned}
$$

If $K < \alpha^{1/d}$, then the term above is strictly positive when $\eta$ is sufficiently close to $\alpha$ and $\kappa$ is sufficiently large. The rest of the proof is analogous to the proof in Section B.4.

## B.8 Proof of Proposition 1

I first formally construct the initial model and the $q$-family. The agent is unsure about the asset-specific variations and update his beliefs over the parameter space $\Omega^\theta = \prod_{n=1}^{N} \Omega_n^\theta \subseteq [\underline{\omega}, \overline{\omega}]^N$. Let $q(\cdot|b, \omega_1, ..., \omega_N, a) \in \Delta\mathcal{Y}$ represent the outcome distribution implied by $\theta$ and $(\omega_n)$ when $\hat{b} = b$; in addition, let $\Omega^q = [\underline{b}, \overline{b}] \times [\underline{\omega}, \overline{\omega}]^N$. Then $\theta$ belongs to the $q$-family if we make the market factor $b$ an additional parameter and write $\overline{\Omega}^\theta = \{\hat{b}\} \times \Omega^\theta$.

If $\hat{b} < \underline{\beta}$, investing in the safe asset is a strictly optimal action and a uniformly strict self-confirming equilibrium. It is self-confirming because the agent has correct beliefs about the constant return $G$. The observation then follows from Corollary 2.

If $\hat{b} > \overline{\beta}$, then investing in the safe asset is strictly suboptimal. However, no matter which risky asset the agent invests, his expected average return is higher than the actual average return because $g(\hat{b}, \underline{\omega}) > g^* = \max_{1 \leq n \leq N} g(b^*, \omega_n^*)$. Hence, there exists no SCE under $\theta$ and it is not globally robust. Moreover, since $f$ is normal, in any BN-E $\sigma$ under $\theta$, the agent's belief assigns probability 1 to $\Omega^\theta(\sigma) = \{\omega \in \Omega^\theta : \omega_n = \min \Omega_n^\theta, \forall n \in \text{supp}(\sigma)\}$. Now consider any competing model $\theta^c$ with $\tilde{b} \in [b^*, \hat{b})$. When $\tilde{b}$ is sufficiently close to $\hat{b}$, we have $g(b^*, \omega_n^*) < g(\tilde{b}, \omega_n) < g(\hat{b}, \omega_n)$ for all $n$ and all $\omega_n \in \Omega_n^\theta$. It follows that

$$\begin{aligned}\mathbb{E}\left(\ln \frac{q(\cdot|\tilde{b}, \omega, a)}{q(\cdot|\hat{b}, \omega, a)}\right) &= \int_{\mathcal{Y}} q(y|b^*, \omega^*, a) \ln \frac{q(y|\tilde{b}, \omega, a)}{q(y|\hat{b}, \omega, a)} dy \\ &= \int_{\mathcal{Y}} f(y - g(b^*, \omega_n^*)) \ln \frac{f(y - g(\tilde{b}, \omega_n))}{f(y - g(\hat{b}, \omega_n))} dy > 0\end{aligned}$$

for all $a = n \in \mathcal{A} \setminus \{N+1\}$. Since $\tilde{b}$ could be arbitrarily close to $\hat{b}$, by Theorem 5, model $\theta$ is not $q$-constrained locally robust. Now, let's maintain that $\tilde{b} \in [b^*, \hat{b})$ but relax the assumption that $\tilde{b}$ is close to $\hat{b}$. Then for all $n \in \{1, ..., N\}$, either $g(b^*, \omega_n^*) < g(\tilde{b}, \min \Omega_n^\theta) < g(\hat{b}, \min \Omega_n^\theta)$, or there exists $\tilde{\omega}_n \in [\underline{\omega}, \overline{\omega}]$ such that $g(\tilde{b}, \tilde{\omega}_n) = g(b^*, \omega_n^*)$, and thus $g(b^*, \omega_n^*) \leq g(\tilde{b}, \tilde{\omega}_n) < g(\hat{b}, \underline{\omega})$. Model $\theta$ does not persist against model $\theta^c$ with $\tilde{\omega}_n \in \Omega_n^{\theta^c}$.

If $\hat{b} \in [\underline{\beta}, \overline{\beta}]$, then model $\theta$ is globally robust because any strategy $\sigma$ that takes support over all assets with their average return equal to $y^*$ is a quasi-strict self-confirming equilibrium.

## B.9 Proof of Proposition 2

Suppose the agent's action space contains $K$ elements, $a^1 < a^2 < ... < a^K$. Define function $h : [\underline{\omega}, \overline{\omega}] \to [\underline{\omega}, \overline{\omega}]$, such that $h(\omega)$ returns the KL minimizer evaluated at the largest myopically optimal action against the degenerate belief $\delta_\omega$ i.e. $h(\omega)$ minimizes $D_{KL}\left(q^*(\cdot|\overline{a}(\omega)) \parallel q(\cdot|\overline{a}(\omega), \hat{b}, \omega)\right)$ where $\overline{a}(\omega) = \max A^\theta(\delta_\omega)$. By Assumption 4, there exists an increasing sequence of intervals $\{(\omega_k, \omega_{k+1})\}_{k=0}^K$ such that $\omega_0 = \underline{\omega}$, $\omega_K = \overline{\omega}$, $a^k$ is the unique myopically optimal action over $(\omega_{k-1}, \omega_k)$ and both $a^{k-1}$ and $a_k$ are myopically optimal at $\omega_{k-1}$. Function $h$ is flat within each interval. If there exists a pure BN-E under model $\theta$, then it must be supported by a degenerate belief at $\omega$ such that $h(\omega) = \omega$. By Assumption 4, any pure BN-E must also be self-confirming, and any mixed BN-E cannot be self-confirming.

Suppose $\hat{b} > b^*$, then $h$ jumps up discontinuously at all cutoffs $\{\omega_k\}_{1 \leq k \leq K-1}$. Suppose there exists no solution to $h(\omega) = \omega$. Then since $h(\underline{\omega}) \geq \underline{\omega}$ and $h(\overline{\omega}) \leq \overline{\omega}$, we know that there must exist $\hat{k}$ such that $h(\omega) > \omega$ for all $\omega \in (\omega_{k^*-1}, \omega_{k^*})$ and $h(\omega') < \omega'$ for all $\omega' \in (\omega_{k^*}, \omega_{k^*+1})$. But this contradicts the observation that $h$ jumps up at $\omega_{k^*}$. It also immediately follows that there exists a solution $\hat{\omega}$ to $h(\hat{\omega}) = \hat{\omega}$ such that $h(\omega') > \omega'$ for $\omega' < \hat{\omega}$ and $h(\omega'') < \omega''$ for $\omega'' < \hat{\omega}$. Let $\hat{a}$ be the unique myopically optimal action at $\delta_{\hat{\omega}}$. Then $\hat{a}$ is a pure self-confirming equilibrium, supported by the generate belief at $\hat{\omega}$. By assumption 4, $\omega \in \Omega^\theta$, and thus $\hat{a}$ is also a self-confirming equilibrium under $\theta$. Note that $\hat{a}$ is uniformly strict. By Corollary 2, model $\theta$ is globally robust.

Now suppose the agent is underconfident, then $h$ jumps down discontinuously at the cutoffs $\{\omega_k\}_{1 \le k \le K-1}$. Hence, there exists at most one solution to $h(\omega) = \omega$. Suppose there exists a SCE $\sigma^\dagger$ when the agent believes his ability is given by $\tilde{b}$. Then by the upper-hemicontinuity of $A^\theta$, when $\hat{b}$ is lower than but sufficiently close to $\tilde{b}$, there exists some $\hat{\omega} > \omega^*$ such that $g(a^\dagger, \hat{b}, \hat{\omega}) = g(a^\dagger, b^*, \omega^*)$, where $a^\dagger = \max \operatorname{supp}(\sigma^\dagger)$ and is the unique myopically optimal action against $\delta_{\hat{\omega}}$. It follows that $a^\dagger$ is a uniformly strict SCE under $\theta$. Since there always exists a SCE when the agent is correctly specified, i.e. $\tilde{b} = b^*$, we infer that model $\theta$ is globally robust when $b^* - \hat{b}$ is sufficiently small.

Suppose instead that there is no solution to $h(\omega) = \omega$ when the agent's self-perception is given by $\hat{b}$. If so, there exists no SCE under model $\theta$. By Theorem 1, $\theta$ is not globally robust. By continuity, $h(\omega) = h(\omega)$ also does not admit any solution at $\tilde{b}$ if it is sufficiently close to $\tilde{b}$. Therefore, there exists an open neighborhood around $\hat{b}$ such that model $\theta$ is not globally robust. I now show that $\theta$ is not $q$-constrained locally robust. Suppose $\theta$ admits a mixed BN-E $\hat{\sigma}$, supported by a potentially mixed belief $\hat{\pi}^\theta \in \Delta\Omega^\theta$. Suppose $\hat{\sigma}$ takes support over $a^k$ and $a^{k+1}$ (note that both $\hat{\sigma}$ and $\hat{\pi}^\theta$ have at most two elements in their support). Then for any $\hat{\omega} \in \operatorname{supp}(\pi^\theta)$, we have

$$\hat{\omega} \in \arg\min_{\Omega^\theta} \sigma(a^k) D_{KL}\left(q^*(\cdot|a^k) \parallel q^\theta(\cdot|a^k, \omega')\right) + \sigma(a^{k+1}) D_{KL}\left(q^*(\cdot|a^{k+1}) \parallel q^\theta(\cdot|a^{k+1}, \omega')\right).$$

Let $\omega^k$ denote the KL minimizer in $[\underline{\omega}, \overline{\omega}]$ at $a^k$ and $\omega^{k+1}$ denote the KL minimizer in $[\underline{\omega}, \overline{\omega}]$ at $a^{k+1}$. Since $g(a, \hat{b}, \omega) - g(a, b^*, \omega^*)$ is strictly increasing in $a$ when $\hat{b} < b^*$ and $\omega > \omega^*$, we know that $\omega^k > \omega^{k+1}$. Since $\omega^k, \omega^{k+1} \in \Omega^\theta$ by assumption, we have $\operatorname{supp}(\hat{\pi}^\theta) \subseteq [\omega^{k+1}, \omega^k]$. Hence, for all $\hat{\omega} \in \operatorname{supp}(\hat{\pi}^\theta)$,

$$g(a^k, \hat{b}, \hat{\omega}) - g(a^k, b^*, \omega^*) \le 0$$
$$g(a^{k+1}, \hat{b}, \hat{\omega}) - g(a^{k+1}, b^*, \omega^*) \ge 0,$$

with at least one inequality being strict. Suppose the second inequality is strict. Pick $\tilde{b} \in (\hat{b}, b^*]$ and define $\tilde{\omega}$ by $g(a^k, \hat{b}, \hat{\omega}) = g(a^k, \tilde{b}, \tilde{\omega})$. Similarly, since $g(a, \hat{b}, \omega) - g(a, \tilde{b}, \tilde{\omega})$ is strictly increasing in $a$, we know that $g(a^{k+1}, \hat{b}, \hat{\omega}) > g(a^{k+1}, \tilde{b}, \tilde{\omega})$; analogously $g(a^{k+1}, \hat{b}, \hat{\omega}) - g(a^{k+1}, b^*, \omega^*) \ge 0$. Therefore, the $\hat{\sigma}$-weighted KL divergence is smaller at $(\tilde{b}, \tilde{\omega})$ than at $(\hat{b}, \hat{\omega})$. When $\tilde{b}$ is sufficiently close to $\hat{b}$, the parameter pair $(\tilde{b}, \tilde{\omega})$ is also close to $(\hat{b}, \hat{\omega})$. Since the agent's action frequency converges, by Theorem 5, model $\theta$ is not $q$-constrained locally robust at $\hat{b}$. Moreover, for any $\tilde{b} \in (\hat{b}, b^*]$, there exists a competing model $\theta^c$ with $(\tilde{b}, \tilde{\omega}) \in \Omega^q$ such that $\theta$ does not persist against $\theta^c$.

Similarly, model $\theta$ is not $q$-constrained locally robust if we perturb the value of $\hat{b}$. Combining this with the previous observation, we could find a sequence of intervals such that model $\theta$ is either globally robust or not $q$-constrained locally robust, each occuring within disjoint intervals.

# C Supplemental Appendix

**Example 5** (A p-absorbing mixed SCE)**.** Suppose there are two actions, $\{1, 2\}$ and three parameters $\{1, 1.5, 2\}$ inside the parameter space of model $\theta$. The agent's payoff is simply the outcome $y_t$, with the true DGP being normal distribution $N(0.25, 1)$. However, model $\theta$ is misspecified and predicts that $y_t \sim N((a_t - \omega)^2, 1)$. Note that every (mixed) action is a self-confirming equilibrium, with the supporting belief assigning probability 1 to the parameter value of 1.5. Here, every mixed SCE is p-absorbing since its support contains every action that can be played by the agent. But her action sequence may never converge. To see that, notice that a belief that assigns larger probability to $\omega = 1$ than $\omega = 2$ leads the agent to play $a = 2$, but such play in turn leads her posterior to drift away from $\omega = 1$ towards $\omega = 2$.

I provide two examples below to substantiate the observation in Footnote 14. Example 6 presents a scenario in which $\theta$ persists against $\theta^1$ and $\theta^2$ separately but does not persist against $\{\theta^1, \theta^2\}$, while Example 7 shows an opposite scenario.

**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 $(x^1, x^2, x^3, x^4)$. 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 $a = 2$. That is, the outcome $y$ is given by $(x^1, x^3)$ when $a = 1$ and given by $(x^2, x^4)$ when $a = 2$. She entertains an initial model $\theta$ 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. The predicted means are independent of the actions taken.

| $\theta$ | $\omega^1$ | $\omega^2$ |
|---|---|---|
| $(\overline{x}^1, \overline{x}^2, \overline{x}^3, \overline{x}^4)$ | $(1, 1, 1, 0)$ | $(1, 1, 0, 1)$ |

| $\theta^1$ | $\omega^{1\prime}$ | $\omega^{2\prime}$ |
|---|---|---|
| $(\overline{x}^1, \overline{x}^2, \overline{x}^3, \overline{x}^4)$ | $(1, 0, 1, 0)$ | $(1, 0, 0, 1)$ |

| $\theta^2$ | $\omega^{1\prime\prime}$ | $\omega^{2\prime\prime}$ |
|---|---|---|
| $(\overline{x}^1, \overline{x}^2, \overline{x}^3, \overline{x}^4)$ | $(0,1,1,0)$ | $(0,1,0,1)$ |

Notice that there are two uniformly strict and thus p-absorbing Berk-Nash equilibria under $\theta$: (1) $a = 1$ is played w.p. 1, supported by the belief that assigns probability 1 to $\omega^1$; (2) $a = 2$ is played w.p. 1 , supported by the belief that assigns probability 1 to $\omega^2$. First observe that $\theta$ persists against $\theta^1$ at a prior $\pi_0^\theta$ that assigns sufficiently high belief to $\omega^1$. This follows from the fact that the likelihood ratio between $\theta$ and $\theta^1$ is always 1 when $a = 1$ is played, and that the equilibrium is p-absorbing. Analogously, $\theta$ persists against $\theta^2$ at a prior $\pi_0^\theta$ that assigns sufficiently high belief to $\omega^2$. However, notice that $\theta$ 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 $\theta^1$ and $\theta^2$ would fit the data strictly better than $\theta$, prompting the agent to adopt $\theta^1$ and $\theta^2$ infinitely often.

**Example 7.** Let $y$ be a normally distributed variable with mean 0 and variance 1, whose distribution is independent of actions. 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 $y$. Her flow payoff is given by $a \cdot y$. She entertains an initial model $\theta$ and two competing models, $\{\theta^1, \theta^2\}$. Model $\theta^1$ has a single parameter and perfectly matches the true DGP, while models $\theta$ and $\theta^2$ both have a binary parameter space. The predictions about $\overline{y}$ of each model are summarized by the following table.

| $\theta$ | $\omega^1$ | $\omega^2$ |
|---|---|---|
| $a^1$ | $-1$ | $1$ |
| $a^2$ | $-2$ | $1$ |

| $\theta^1$ | $\omega'$ |
|---|---|
| $a^1, a^2$ | $-1$ |

| $\theta^2$ | $\omega''$ |
|---|---|
| $a^1, a^2$ | $2$ |

Suppose the agent's prior satisfies that $\pi_0^\theta(\omega^1) = 1 - \pi_0^\theta(\omega^0) = \frac{0.5}{\alpha} < \frac{1}{\alpha}$. First consider what happens when the agent has only one competing model, $\theta^1$. By the Law of Large

Numbers, the likelihood ratio between $\theta^1$ and $\theta$ eventually exceeds $\alpha$ almost surely because

$$\begin{aligned}
\frac{l_t^{\theta^1}}{l_t^{\theta}} &= \frac{\prod_{\tau=0}^{t-1} q^{\theta^1}\left(y_\tau|a_\tau,\omega'\right)}{\prod_{\tau=0}^{t-1} q^{\theta}\left(y_\tau|a_\tau,\omega^1\right)\pi_0^{\theta}\left(\omega^1\right)+\prod_{\tau=0}^{t-1} q^{\theta}\left(y_\tau|a_\tau,\omega^2\right)\pi_0^{\theta}\left(\omega^2\right)} \\
&= \frac{\prod_{\tau=0}^{t-1} q^{*}\left(y_\tau|a_\tau\right)}{\prod_{\tau=0}^{t-1} \mathbf{1}_{(a_\tau=a^1)} q^{*}\left(y_\tau|a_\tau\right)\pi_0^{\theta}\left(\omega^1\right)+\xi(h_t)} \\
&\geq \frac{\prod_{\tau=0}^{t-1} q^{*}\left(y_\tau|a_\tau\right)}{\prod_{\tau=0}^{t-1} \frac{1}{\alpha} q^{*}\left(y_\tau|a_\tau\right)+\xi(h_t)}
\end{aligned}$$

where $\frac{\xi(h_t)}{\prod_{\tau=0}^{t-1} q^*(y_\tau|a_\tau)}$ converges to 0 almost surely. Threrefore, $\theta$ does not persist against $\theta^1$ under prior $\pi_0^{\theta}$.

However, model $\theta$ persists against $\Theta^c := \{\theta^1, \theta^2\}$ at prior $\pi_0^{\theta}$. First notice that for any $a_0 \in \mathcal{A}$, there exists some $y_0$ sufficiently large such that

$$l_1^{\theta^2} > \alpha \cdot \max\{l_1^{\theta}, l_1^{\theta^1}\}$$

and thus the agent switches to $\theta^2$ in the beginning of period 1. As a result, the agent plays $a_1 = a^2$ in period 1 since it is the strictly dominant strategy under $\theta^2$. But then we could find some sufficiently small $y_1$ such that the following two inequalities hold:

$$l_2^{\theta} > \alpha \cdot \max\{l_2^{\theta^1}, l_2^{\theta^2}\},$$

$$\pi_2^{\theta}(\omega^1) = \frac{\pi_0^{\theta}(\omega^1) q^{\theta}(y_0|a_0,\omega^1) q^{\theta}(y_1|a_1,\omega^1)}{\sum_{\omega\in\{\omega^1,\omega^2\}} \pi_0^{\theta}(\omega) q^{\theta}(y_0|a_0,\omega) q^{\theta}(y_1|a_1,\omega)} > \max\{\frac{1}{\alpha}, c\},$$

where $c$ is chosen such that while adopting $\theta^1$, the agent finds $a^1$ payoff-maximizing if her belief assigns a probability higher than $c$ to $\omega^1$, i.e. $\pi^{\theta}(\omega^1) > c$. The first inequality implies that the agent switches back to $\theta$ in the beginning of period 2. The second inequality, together with the observation that the pure strategy $a^1$ is a uniformly strict self-confirming equilibrium supported by the belief that assigns probability 1 to $\omega^1$, ensures that with positive probability, the agent plays action $a^1$ forever, provided that her decisions are made based on model $\theta$. But notice that on those paths, the agent indeed no longer switches to other models after period 1 because for $t > 2$,

$$\frac{l_t^{\theta^1}}{l_t^{\theta}} < \frac{l_2^{\theta^1}}{l_2^{\theta}} \frac{\prod_{\tau=2}^{t-1} q^{*}\left(y_\tau|a_\tau\right)}{\prod_{\tau=2}^{t-1} q^{*}\left(y_\tau|a_\tau\right)\pi_2^{\theta}\left(\omega^1\right)} < 1 < \alpha.$$

Since outcomes $y_0$ and $y_1$ that satisfy the aformentioned properties are drawn with positive probability, we conclude that $\theta$ persists against $\Theta^c := \{\theta^1, \theta^2\}$ at prior $\pi_0^{\theta}$.