

Evolutionarily Stable (Mis)specifications: Theory and Applications*

Kevin He[†] Jonathan Libgober[‡]

First version: December 20, 2020

This version: May 9, 2021

Abstract

We introduce an evolutionary framework to evaluate competing (mis)specifications in strategic situations, focusing on which misspecifications can persist over a correct specification. Agents with heterogeneous specifications coexist in a society and repeatedly match against random opponents to play a stage game. They draw Bayesian inferences about the environment based on personal experience, so their learning depends on the distribution of specifications and matching assortativity in the society. One specification is *evolutionarily stable* against another if, whenever sufficiently prevalent, its adherents obtain higher expected objective payoffs than their counterparts. The learning channel leads to novel stability phenomena compared to frameworks where the heritable unit of cultural transmission is a single belief instead of a specification (i.e., set of feasible beliefs). We apply the framework to linear-quadratic-normal games where players receive correlated signals but possibly misperceive the information structure. The correct specification is not evolutionarily stable against a correlational error, whose direction depends on matching assortativity. As another application, the framework also endogenizes coarse analogy classes in centipede games.

Keywords: misspecified Bayesian learning, endogenous misspecifications, evolutionary stability, higher-order beliefs, analogy classes

*We thank Cuimin Ba, In-Koo Cho, Krishna Dasaratha, Andrew Ellis, Ignacio Esponda, Mira Frick, Drew Fudenberg, Alice Gindin, Ryota Iijima, Yuhta Ishii, Philippe Jehiel, Pablo Kurlat, Filippo Massari, Andy Postlewaite, Philipp Sadowski, Alvaro Sandroni, Joshua Schwartzstein, Philipp Strack, and various conference and seminar participants for helpful comments. Kevin He thanks the California Institute of Technology for hospitality when some of the work on this paper was completed.

[†]University of Pennsylvania. Email: hesichao@gmail.com

[‡]University of Southern California. Email: libgober@usc.edu

1 Introduction

In many economic settings, people draw *misspecified* inferences about the world — that is, they start with a prior belief that dogmatically precludes the true data-generating process. For instance, behavioral economics documents a number of prevalent statistical biases. When people reason about economic fundamentals under the spell of one of these biases, they engage in misspecified learning. Following [Esponda and Pouzo \(2016\)](#), a growing literature has focused on the implications of Bayesian learning under different misspecifications, taking the errors as exogenously given.

Why and when might we expect such misspecifications to persist? Mistakes that distort learning are empirically ubiquitous, which is puzzling for two reasons. First, many of these errors demand even greater computational sophistication than the simple truth, making them hard to justify on the grounds of bounded cognition or costly attention. Convoluted conspiracy theories fall into this category. So does a behavioral error called projection bias, where agents overestimate the similarity between their own information and others' information. Reasoning with projection bias in settings with statistical independence requires the learner to keep track of inter-personal correlations, complicating the inference problem. Second, conventional economic wisdom suggests competitive pressure eliminates mistakes — including misspecifications. Indeed, contemporaneous papers that formalize payoff-based criteria for selecting between (mis)specifications find no strict advantage to deviating from a dogmatically correct view of the data-generating process in single-agent decision problems ([Fudenberg and Lanzani, 2020](#); [Frick, Iijima, and Ishii, 2021](#)).

This paper introduces a general framework to evaluate competing (mis)specifications based on their expected objective payoffs, with particular emphasis on which misspecifications are likely to persist over the correct specification (and in which environments). Agents draw inferences in response to their observations, and their equilibrium beliefs play a critical role in selecting between different specifications and generate novel stability phenomena. We find that when agents with heterogeneous specifications coexist in a society and repeatedly match against random opponents to play a stage game, misspecified agents may enjoy a strict payoff advantage compared to their correctly specified counterparts. Unlike in decision problems, misspecifications in games can lead to *strategically* beneficial misinferences about the game parameters. Through several examples and applications, we discuss how details of the social interaction structure, such as the matching assortativity between agents with different specifications, shape the stability of different mistakes.

1.1 Correlational Mistakes in a Linear-Quadratic-Normal Game

We informally describe an application of our framework. We fix a linear-quadratic-normal game of incomplete information from [Vives \(1988\)](#) as the stage game. There is a population of players who match in pairs every period to play the stage game. Nature’s type is drawn i.i.d. across games, and every pair of matched players receive correlated information about Nature’s type in their game (as in [Bergemann and Morris \(2013\)](#)). For concreteness, think of the players as competing firms and Nature’s type as a demand state. Firms privately observe a signal about today’s demand shock before choosing how much to produce.

The population initially consists of a homogeneous group of firms who have correct beliefs about all game parameters, including the correlation between rival players’ private signals. But now a small fraction of new firms enter the society. The entrants differ from the incumbents in two ways. First, they hold a dogmatically wrong belief about the signal correlation. Second, they are uncertain about a parameter of the stage game — the elasticity of market price with respect to total supply — and learn this fundamental from the realized prices in their games across different periods. The entrants therefore engage in misspecified learning: after seeing their own signals, they hold wrong beliefs about rivals’ signals and hence rivals’ production, so they misinterpret the market price and mislearn game parameters.

Will the market drive out the mistaken entrants? The answer depends on the nature of the mistake and the interaction structure in the society — specifically, the matching assortativity of how incumbents and entrants are paired with each other. Suppose matching is uniform. If the misspecified entrants are slightly biased in the direction of believing in excessively correlated information (projection bias), then they will end up with objectively higher profits than the correctly specified incumbents and grow in relative prominence. On the other hand, if the entrants instead believe in excessively independent information (correlation neglect), they will underperform compared to the incumbents and get driven out. But when matching is perfectly assortative between the incumbent and entrant firms, the conclusion is reversed. In this environment, it is the mistake of correlation neglect that can invade a rational society.

We also use this game to illustrate that the mislearning channel is crucial to the predictions. The correlational errors that persist with learning would instead confer an evolutionary disadvantage if they were combined with correct beliefs about the price elasticity parameter.

1.2 A Framework of Competing Specifications

More generally, we propose an evolutionary framework where specifications are encoded in *theories* that delineate feasible beliefs and serve as the basic unit of cultural transmission. Each theory may represent, for example, a scientific paradigm that stipulates a set of (possibly incorrect) relationships between environmental parameters and observables. The theory’s

adherents estimate its parameters and play the stage game based on their calibrated model. Theories rise and fall in prominence based on the objective welfare of their adherents, as the school of thought that leads to higher payoffs tends to acquire more resources and attract more followers in the future.

In the example above, there are two theories about the signal correlation of rival firms that sell to the same market. Every firm learns about a parameter of the environment (price elasticity) through the lens of their theory. Firms that believe in different correlations interpret the same observation differently when inferring price elasticity, as they make different estimates about rival firm's production based on their own demand signal. We suppose that more future entrepreneurs flock to the theory that leads to objectively higher firm profits.

The fitness of a theory is determined by its average payoff in stage games, and this average depends on the distribution of opponents. We introduce the concept of a *zeitgeist* to capture the relevant social interaction structure in the society — the sizes of the subpopulations with different theories, and the matchmaking technology that pairs up opponents to play the game. In equilibrium, each agent forms a Bayesian belief about her environment using data from all of her interactions, and subjectively best responds to this belief. We define the *evolutionary stability* of theory A against theory B based on whether theory A has a weakly higher equilibrium fitness than theory B when the population share of theory A is close to 1. The example above, for instance, says the correct specification about the information structure of firms is evolutionarily unstable against either projection bias or correlation neglect, depending on the matching assortativity.

Adherents of a misspecified theory may come to different conclusions about the economic fundamentals in different zeitgeists, with these different beliefs translating into different subjective best-response functions in the stage game. This kind of endogeneity in stage-game behavior leads to novel stability phenomena. First, we show the possibility of a strong form of multiplicity in the stability comparison between two theories: *stability reversals*. Two theories exhibit stability reversal if (i) theory A's adherents strictly outperform theory B's adherents not only on average, but even conditional on opponent's type, whenever theory A is dominant; and (ii) theory B's adherents strictly outperform theory A's adherents, whenever theory B is dominant. Second, we show that the relative stability of one theory over another may be non-monotonic in matching assortativity. One theory may be evolutionarily stable against another when assortativity is either high or low, but not when it is intermediate. Both of these stability phenomena operate through misinference and cannot happen if the learning channel is eliminated. That is, they never arise in a world where the basic unit of cultural transmission is a single belief about the economic environment instead of a theory (i.e., a collection of feasible beliefs).

The rest of this section reviews related literature. Section 2 introduces the environ-

ment and the evolutionary framework for assessing the stability of specifications. Section 3 discusses how the learning channel enables novel stability phenomena. Section 4 contains application to misspecified information structures in linear-quadratic-normal games. Section 5 contains a second application that shows our framework can endogenize analogy classes, a solution concept that Jehiel (2005) introduced to capture simplified strategic thinking in complex environments. It also illustrates how payoff structure of the stage game can alter the prevalence of analogy-based thinking in the society. Section 6 concludes. Appendix A contains proofs of results from the main text. Appendix B presents sufficient conditions for the existence and upper hemicontinuity of equilibrium zeitgeists, and Appendix C provides a learning foundation for equilibrium zeitgeists.

1.3 Related Literature

Our paper contributes to the literature on misspecified Bayesian learning by proposing a framework to assess which specifications are more likely to persist based on their objective performance. Most prior work on misspecified Bayesian learning study implications of particular errors in specific active-learning environments (i.e., when actions affect observations), including both single-agent decision problems (Nyarko, 1991; Fudenberg, Romanyuk, and Strack, 2017; Heidhues, Koszegi, and Strack, 2018; He, 2020) and multi-agent games (Bohren, 2016; Bohren and Hauser, 2018; Jehiel, 2018; Molavi, 2019; Dasaratha and He, 2020; Ba and Gindin, 2020; Frick, Iijima, and Ishii, 2020; Murooka and Yamamoto, 2021). A number of papers establish general convergence properties of misspecified learning (Esponda and Pouzo, 2016; Esponda, Pouzo, and Yamamoto, 2019; Frick, Iijima, and Ishii, 2019; Fudenberg, Lanzani, and Strack, 2020). All of the above papers take misspecifications as exogenously given. By contrast, we propose endogenizing misspecifications using ideas from evolutionary game theory. This also lets us ask how details of the evolutionary process (e.g., the matching assortativity) shape the stability of misspecifications.

Another strand of literature shares our central focus on selecting between multiple specifications for Bayesian learning. Papers in this literature have focused on different criteria, including performance in financial markets (Sandroni, 2000; Massari, 2020), subjective expectations of payoffs (Olea, Ortoleva, Pai, and Prat, 2020; Levy, Razin, and Young, 2020; Eliaz and Spiegel, 2020; Gagnon-Bartsch, Rabin, and Schwartzstein, 2020), and goodness-of-fit tests (Cho and Kasa, 2015, 2017; Ba, 2020; Schwartzstein and Sunderam, 2021). We instead consider the objective expected payoffs of agents with different specifications who coexist in the same society and interact strategically. We are implicitly motivated by a story of cultural transmission where agents play “games” with random opponents and derive welfare based on these interactions, and those with higher welfare are more likely to pass down

their theories to future agents.

This paper is closest to two independent and contemporaneous papers, [Fudenberg and Lanzani \(2020\)](#) and [Frick, Iijima, and Ishii \(2021\)](#), who consider welfare-based criteria for selecting among misspecifications in single-agent decision problems. [Fudenberg and Lanzani \(2020\)](#) study a framework where a continuum of agents with heterogeneous misspecifications arrive each period and learn from their predecessors’ data. When the population shares of different misspecifications change according to their objective performance, [Fudenberg and Lanzani \(2020\)](#) ask which Berk-Nash equilibria under one misspecification are robust to invasion by a small fraction of mutants with a different misspecification. [Frick, Iijima, and Ishii \(2021\)](#) assign a *learning efficiency index* to every misspecified signal structure and conduct a robust comparison of welfare under different misspecifications. For two misspecifications with the property that biased agents still learn the correct state in the long run, the misspecification with a higher index leads to faster convergence to the truth and thus higher welfare in any decision problem, provided there is a large enough but finite number of signals.

In single-agent decision problems, correctly specified agents always perform weakly better than misspecified agents in the long run (except when there are non-identifiability issues, see [Proposition 1](#)), so welfare-based criteria do not provide a strict advantage in equilibrium to misspecified individuals compared to the correctly specified ones in the same society.¹ [Frick, Iijima, and Ishii \(2021\)](#) also find that correctly specified agents who know the data-generating process converge to the truth faster than misspecified agents.² By contrast, we focus on a theory of welfare-based selection of misspecifications in games, where strategic concerns may imply that learning under a misspecification confers a strict evolutionary advantage relative to learning under the correct specification.³ The central concept in our framework, a *zeitgeist*, captures aspects of the social interaction structure that are uniquely relevant when agents confront a game as opposed to a decision problem — namely, the assortativity of the matching technology that pairs up agents with different specifications to play the stage game, and how agents behave when matched against different types of opponents.

Our framework of competition between different specifications for Bayesian learning is

¹This conclusion relates to the market-selection hypothesis that dates back to at least [Friedman \(1953\)](#). [Blume and Easley \(2002, 2006\)](#) come to the same conclusion in market equilibria where agents act as price-takers, provided the market environment leads to Pareto efficient outcomes. Our environment where agents pair off into two-player games and interact as local oligopolies is a natural setting where outcomes are not Pareto efficient.

²But, correctly specified agents who are uncertain about the data-generating process may converge more slowly than misspecified agents.

³Some papers studying misspecified learning in games also point out that misspecifications can improve an agent’s welfare in particular situations (e.g., [Jehiel \(2005\)](#) and [Ba and Gindin \(2020\)](#)). We contribute by introducing a general framework that can be applied broadly.

inspired by the evolutionary game theory literature. This literature also uses objective payoffs as the selection criterion, and studies the evolution of subjective preferences in games and decision problems (e.g., Dekel, Ely, and Yilankaya (2007), see also the surveys Robson and Samuelson (2011) and Alger and Weibull (2019)) and the evolution of constrained strategy spaces (Heller, 2015; Heller and Winter, 2016). Learning does not play a key role in these papers. By contrast, our work seeks to provide a foundation for the exogenously given misspecifications in the recent literature on misspecified Bayesian learning, and our results depend crucially on the mislearning channel (as highlighted in Section 3 and Section 4.6). In settings where agents entertain fundamental uncertainty about payoff parameters in the stage game, we can think of our framework as applying evolutionary forces to *sets of preferences*, viewing every misspecification (i.e., a set of feasible stage-game parameters) as a set of preferences over strategy profiles. A few papers in this literature study the evolution of different belief-formation processes (Heller and Winter, 2020; Berman and Heller, 2020), but they take a reduced-form (and possibly non-Bayesian) approach and consider arbitrary inference rules. We require agents to be Bayesians who only differ in the support of their Bayesian prior (i.e., their specification), given the relation of this work to the literature on misspecified Bayesian learning.

2 Environment and Stability Concept

In this section, we introduce the general environment and stability concept. We begin with the objective stage game and subjective theories that encode specifications. We define the notion of an *equilibrium zeitgeist*, which describes the steady-state behavior and beliefs in a society populated by agents with heterogeneous specifications. We then present the stability concept, based on objective welfare in equilibrium zeitgeists when one theory is sufficiently prevalent. We also discuss an extension, *equilibrium zeitgeist with strategic uncertainty*, where agents may hold wrong beliefs about others' strategies in the steady state.

2.1 Objective Primitives

We first set up the objective primitives of the general environment. The stage game is a symmetric two-player game with a common strategy space \mathbb{A} , assumed to be metrizable. When i and $-i$ choose strategies $a_i, a_{-i} \in \mathbb{A}$, random consequences $y_i, y_{-i} \in \mathbb{Y}$ are generated for the players from a metrizable space \mathbb{Y} . These consequences determine each player's utility, according to a utility function $\pi : \mathbb{Y} \rightarrow \mathbb{R}$. Objectively, y_i is generated as a function of i and $-i$'s play. We take this distribution to be $F^\bullet(a_i, a_{-i}) \in \Delta(\mathbb{Y})$, where $\Delta(\mathbb{Y})$ is the set of distributions over \mathbb{Y} . We denote the density or probability mass function associated

with this distribution by $f^\bullet(a_i, a_{-i}) : \mathbb{Y} \rightarrow \mathbb{R}_+$.

This general setup can allow for mixed strategies (if \mathbb{A} is the set of mixtures over some pure actions) and incomplete-information games (if S is a space of private signals, A a space of actions, and $\mathbb{A} = A^S$ is the set of signal-contingent actions). It can also describe asymmetric games. Suppose there is a game with action sets A_1, A_2 for player roles P1 and P2, and that the consequences of P1 and P2 under the action profile $(a_1, a_2) \in A_1 \times A_2$ are generated according to the distributions $F_1^\bullet(a_1, a_2)$ and $F_2^\bullet(a_2, a_1)$ over \mathbb{Y} , where we assume the consequence also fully reveals the agent's role. We may construct a symmetric stage game by letting $\mathbb{A} = A_1 \times A_2$, so the strategies of two matches agents spell out what actions they would take if they were assigned into each of the player roles. The agents are then placed into the player roles uniformly at random and play according to the strategies. That is, the objective distribution over i 's consequence when playing $(a_{i1}, a_{i2}) \in \mathbb{A}$ against $(a_{-i1}, a_{-i2}) \in \mathbb{A}$ is given by the 50-50 mixture over $F_1^\bullet(a_{i1}, a_{-i2})$ and $F_2^\bullet(a_{i2}, a_{-i1})$.

2.2 Models and Theories

Throughout this paper, we will take the strategy space \mathbb{A} , the set of consequences \mathbb{Y} , and the utility function over consequences π to be common knowledge among the agents. But, agents are unsure about how play in the stage game translates into consequences — that is, they have *fundamental uncertainty* about the function F^\bullet . For example, the agents may be uncertain about some parameters of the stage game, such as the market price elasticity in a quantity-competition game.

We will consider a society with two observably distinguishable groups of agents, A and B, who may behave differently in the stage game (due to each group having a different belief about the economic fundamentals, for example). All agents entertain different *models* of the world as possible resolutions of their uncertainty. A model $F : \mathbb{A}^2 \rightarrow \Delta(\mathbb{Y})$ is a conjecture about how strategy profiles translate into consequences for the agent. Assume each F , like F^\bullet , is given by a density or probability mass function $f(a_i, a_{-i}) : \mathbb{Y} \rightarrow \mathbb{R}_+$ for every $(a_i, a_{-i}) \in \mathbb{A}^2$.

A *theory* Θ is a collection of models: that is, a subset of $(\Delta(\mathbb{Y}))^{\mathbb{A}^2}$. We assume Θ is metrizable. Each agent enters society with a persistent theory, which depends entirely on whether they are from group A or group B. We think of this exogenously endowed theory as coming from education or cultural background, and each agent dogmatically believes that her theory contains the correct model of the world. A theory Θ is *correctly specified* if $F^\bullet \in \Theta$.

In general, a theory may exclude the true fundamental environment F^\bullet . Such *misspecified* theories can represent a scientific paradigm about the economy based on a false premise, a religious belief system with dogmas that contradict facts about the world, or heuristic

thinking stemming from a psychological bias that deems the true environment as implausible. Each agent plays the stage game with a random opponent in every period, and uses her personal experience in these matches to calibrate the most accurate model within her theory in a way that we will make precise in Section 2.4.

An agent endowed with a theory is called an *adherent* of the theory. As alluded to above, we suppose the society is composed of the adherents in the two observable groups A and B. This presumes that agents can identify which group their matched opponent belongs to, though we do not assume that agents know the models contained in theories other than their own. For instance, imagine two major theories of business economics coexist in a society, taught by two different universities. Agents are the executives of competing firms and they can use public records to look up the educational background of other executives and therefore learn which school of thought they subscribe to. But even though each agent can perfectly identify her opponent’s group membership (which helps to predict the opponent’s behavior), she does not understand anything about the contents of the rival economic theory.

2.3 Zeitgeists

To study competition between two theories, we must describe the social composition and interaction structure in the society where learning takes place. We introduce the concepts of zeitgeists and equilibrium zeitgeists to capture these details.

The Cambridge Dictionary defines the noun “zeitgeist” as “the general *set of* ideas, beliefs, feelings, etc. that is typical of a particular period in history.” Crucial in this dictionary definition is the *multiplicity* of coexisting ideas and beliefs in the society at a moment in time. In the spirit of the usual meaning of the word, we define a zeitgeist as a *landscape* of concurrent beliefs from different schools of thought, their relative prominence in the society, and the interaction among the adherents of different theories.

Definition 1. A *zeitgeist* $\mathfrak{Z} = (\Theta_A, \Theta_B, \mu_A, \mu_B, p, \lambda, a)$ consists of: (1) two theories Θ_A and Θ_B ; (2) a belief over models for each theory, $\mu_A \in \Delta(\Theta_A)$ and $\mu_B \in \Delta(\Theta_B)$; (3) relative sizes of the two groups in the society, $p = (p_A, p_B)$ with $p_A, p_B \geq 0$, $p_A + p_B = 1$; (4) a matching assortativity parameter $\lambda \in [0, 1]$; (5) each group’s strategy when matched against each other group, $a = (a_{AA}, a_{AB}, a_{BA}, a_{BB})$ where $a_{g,g'} \in \mathbb{A}$ is the strategy that an adherent of Θ_g plays against an adherent of $\Theta_{g'}$.

A zeitgeist outlines the beliefs and interactions among agents with heterogeneous theories living in the same society. Parts (1) and (2) of this definition capture the beliefs of each group. Parts (3) and (4) determine social composition and social interaction—the relative prominence of each theory and the probability of interacting with one’s own group versus with

the population as a whole. In each period, every agent is matched with an opponent from her own group with probability λ , and matched uniformly by population proportion with probability $1 - \lambda$. Therefore, an agent from group g has an overall probability of $\lambda + (1 - \lambda)p_g$ of being matched with an opponent from her own group, and a complementary chance of being matched with an opponent from the other group. Part (5) describes behavior in the society.

2.4 Equilibrium Zeitgeists

To determine each theory’s evolutionary fitness, we introduce our equilibrium concept.

An equilibrium zeitgeist (EZ) imposes equilibrium conditions on the beliefs and behavior in a zeitgeist. Specifically, it is a zeitgeist that satisfies the optimality of inference and behavior, holding fixed the population shares p and the matching assortativity λ . Optimality of behavior requires each player to best respond given her beliefs, and optimality of inference requires that the support of each player’s belief only contains the “best-fitting” models from her theory in the sense of minimizing Kullback-Leibler (KL) divergence.

EZs have a learning foundation (Appendix C) as the social steady state when a continuum of long-lived Bayesian learners with different theories coexist in the society with proportions p , match up with assortativity λ every period, and learn from personal experience. In the learning foundation, each agent starts with a full-support prior belief over the models in her theory and over how others play.⁴ When matched with an opponent, the agent sees the opponent’s group and chooses a strategy $a_i \in \mathbb{A}$. At the end of the game, the agent observes a consequence $y_i \in \mathbb{Y}$ and an ex-post signal x_i about the matched opponent’s strategy a_{-i} . She then updates her belief using Bayes’ rule. When the ex-post signal is sufficiently informative about the matched opponent’s strategy, each agent must hold correct beliefs about others’ strategies in the steady state, and her posterior belief about the fundamental F^\bullet concentrates around those models in her theory that minimize KL divergence.

In an EZ, we abstract away from the issues surrounding learning others’ strategies and focus on learning parameters of the environment. In the example from the introduction where firms match into pairs to play an incomplete-information Cournot duopoly, we can imagine that each firm’s strategy (i.e., mapping from signal realization to production level) corresponds to a manufacturing process that automatically adjusts production based on daily private information, and each firm learns something about the matched rival firm’s production process. We consider an extension in Section 2.7 that allows agents to be misspecified about others’ strategies and hold wrong beliefs about these strategies in equilibrium: this

⁴This setup allows agents to make inferences about game parameters using opponents’ strategy, because their prior may exhibit correlation between strategic uncertainty and fundamental uncertainty.

corresponds to a learning environment where agents only learn from the consequences y_i and do not observe additional feedback about others' strategies.

To formally give the definition of EZ, we require some new notation.

For two distributions over consequences, $\Phi, \Psi \in \Delta(\mathbb{Y})$ with density functions / probability mass functions ϕ, ψ , define the KL divergence from Ψ to Φ as $D_{KL}(\Phi \parallel \Psi) := \int \phi(y) \ln \left(\frac{\phi(y)}{\psi(y)} \right) dy$. Recall that every model F , like the true fundamental F^\bullet , outputs a distribution over consequences for every profile of own play and opponent's play, $(a_i, a_{-i}) \in \mathbb{A}^2$. For model F , let $K(F; a_i, a_{-i}) := D_{KL}(F^\bullet(a_i, a_{-i}) \parallel F(a_i, a_{-i}))$ be the KL divergence from the expected distribution of consequences $F(a_i, a_{-i})$ to the objective distribution of consequences $F^\bullet(a_i, a_{-i})$ under the play (a_i, a_{-i}) .

For a distribution μ over models, let $U_i(a_i, a_{-i}; \mu)$ represent i 's subjective expected utility under the belief that the true model is drawn according to μ . That is, $U_i(a_i, a_{-i}; \mu) := \mathbb{E}_{F \sim \mu}(\mathbb{E}_{y \sim F(a_i, a_{-i})}[\pi(y)])$.

Definition 2. A zeitgeist $\mathfrak{Z} = (\Theta_A, \Theta_B, \mu_A, \mu_B, p, \lambda, a)$ is an *equilibrium zeitgeist (EZ)* if for every $g, g' \in \{A, B\}$, $a_{g,g'} \in \arg \max_{a_i \in \mathbb{A}} U_i(a_i, a_{g',g}; \mu_g)$ and, for every $g \in \{A, B\}$, the belief μ_g is supported on

$$\arg \min_{F \in \Theta_g} \{(\lambda + (1 - \lambda)p_g) \cdot K(F; a_{g,g}, a_{g,g}) + (1 - \lambda)(1 - p_g) \cdot K(F; a_{g,-g}, a_{-g,g})\}$$

where $-g$ means the group other than g .

In an EZ, each agent from group g chooses a subjective best response $a_{g,g'}$ against each group g' of opponents, given her belief μ_g about the fundamental uncertainty. Her belief μ_g is supported on the models in her theory that minimize a weighted KL-divergence objective, with the data from each type of match weighted by the probability of confronting this type of opponent.

In every EZ, agents hold correct beliefs about the strategy of every type of opponent. We focus attention on agents' inferences about the game parameters and abstract away from issues associated with learning how others play. In the learning foundation, we assume the ex-post signal x_i about the matched opponent's strategy is sufficiently accurate, so agents end up with correct beliefs about every group's strategy in the social steady state. (Section 2.7 considers the more general setup where agents entertain both strategic uncertainty and fundamental uncertainty and may hold incorrect beliefs about how others play in equilibrium.)

An important assumption behind this framework is that agents (correctly) believe the economic fundamentals are fixed, no matter who they are matched against. That is, the mapping $(a_i, a_{-i}) \mapsto \Delta(\mathbb{Y})$ describes the stage game that they are playing, and agents know

that they always play the same stage game even though opponents from different groups may use different strategies in the game. As a result, the agent’s experience in games against both groups of opponents jointly resolve the same fundamental uncertainty about the environment. Generally, play between two groups g and g' is not a Berk-Nash equilibrium (Esponda and Pouzo, 2016), as the individuals in group g draw inferences about the game’s parameters not only from the matches against group g' , but also from the matches against the other group $-g'$, who may use a different strategy.

Even as agents adjust their beliefs and behavior to converge to an EZ, the population proportions of different theories p_A, p_B remain fixed. We imagine a world where the relative prominence of theories change much more slowly than the rate of convergence to an EZ. Thus, an equilibrium zeitgeist provides a snapshot of the society in a given era, and the social transitions between different EZs as p evolves takes place on a longer timescale.

2.5 Evolutionary Stability of Theories

In an EZ, define the *fitness* of each theory as the objective expected payoff of its adherents. Consider an evolutionary story where the relative prominence of the two theories in the society rise and fall according to their relative fitness. This could happen, for example, if the theories are the basic heritable units of information passed down to future agents via cultural transmission, and the school of thought whose adherents have higher average payoff tends to acquire more resources and attract a larger share of future adherents. We are interested in a notion of stability based on this “evolutionary” process where two co-existing rival theories compete to create intellectual descendants in a payoff-monotonic way. Can the adherents of a *resident theory* Θ_A , starting at a position of social prominence, always repel an invasion from a small ϵ mass of agents who adhere to a *mutant theory* Θ_B ? The definition of evolutionary stability formalizes this idea.

Since we are motivated by situations where a small but strictly positive population of theory Θ_B adherents invades an otherwise homogeneous society all believing in theory Θ_A , we begin with a refinement of EZ that rules out those equilibria with the population share $(p_A, p_B) = (1, 0)$ that cannot be written as the limit of equilibria with a positive but vanishing p_B . This rules out, for example, EZs with $p_A = 1$ sustained only because group A holds fragile beliefs about the economic fundamentals that would be discarded after a single match against a group B opponent.

Definition 3. An EZ $\mathfrak{Z} = (\Theta_A, \Theta_B, \mu_A, \mu_B, p, \lambda, a)$ with $p = (1, 0)$ is *approachable* if there exists a sequence of EZs $\mathfrak{Z}^{(n)} = (\Theta_A, \Theta_B, \mu_A^{(n)}, \mu_B^{(n)}, (p_A^{(n)}, p_B^{(n)}), \lambda, (a_{AA}^{(n)}, a_{AB}^{(n)}, a_{BA}^{(n)}, a_{BB}^{(n)}))$, where $p_B^{(n)} > 0$ for all n , $p_B^{(n)} \rightarrow 0$, $\mu_A^{(n)} \rightarrow \mu_A$, $\mu_B^{(n)} \rightarrow \mu_B$, $a^{(n)} \rightarrow a$.

In this definition, $\mu_g^{(n)} \rightarrow \mu_g$ refers to convergence in weak* topology on the space $\Delta(\Theta_g)$ of distributions over the models in theory Θ_g , and $a^{(n)} \rightarrow a$ means the convergence of the strategy profile in the metrizable space \mathbb{A}^4 .

We now turn to the definition of evolutionary stability, which is defined only when the set of approachable EZ with $p = (1, 0)$ is non-empty. Stability is defined based on the fitness of theories Θ_A, Θ_B in such equilibria. Evolutionary stability is when Θ_A has higher fitness than Θ_B in all approachable equilibria, and evolutionary fragility is when Θ_A has lower fitness in all approachable equilibria. These two cases give sharp predictions about whether a small share of mutant-theory invaders might grow in size, across all equilibrium selections. A third possible case, where Θ_A has lower fitness than Θ_B in some but not all approachable equilibria, correspond to a situation where the mutant theory may or may not grow in the society, depending on the equilibrium selection.

Definition 4. Suppose there exists at least one approachable EZ with theories Θ_A, Θ_B , $p = (1, 0)$, and matching assortativity λ . Say Θ_A is *evolutionarily stable [fragile]* against Θ_B under λ -matching if in all such approachable EZs, Θ_A has a weakly higher [strictly lower] fitness than Θ_B .

2.6 Misspecified Theories in Decision Problems

In single-agent problems, evolutionary arguments will always favor a correctly specified theory over an incorrect one. The stage “game” is a *decision problem* if $(a_i, a_{-i}) \mapsto F^\bullet(a_i, a_{-i})$ only depends on a_i . In decision problems, the correctly specified theory is evolutionarily stable against any other theory, except when there are identification issues. We adapt the notion of strong identification from [Esponda and Pouzo \(2016\)](#).

Definition 5. Theory Θ_A is *strongly identified* in EZ $\mathfrak{Z} = (\Theta_A, \Theta_B, \mu_A, \mu_B, p, \lambda, a)$ if whenever $F', F'' \in \Theta_A$ both solve

$$\min_{F \in \Theta_A} \{(\lambda + (1 - \lambda)p_A) \cdot K(F; a_{AA}, a_{AA}) + (1 - \lambda)(1 - p_A) \cdot K(F; a_{AB}, a_{BA})\},$$

we have $F'(a_i, a_{AA}) = F''(a_i, a_{AA})$ and $F'(a_i, a_{BA}) = F''(a_i, a_{BA})$ for all $a_i \in \mathbb{A}$.

Proposition 1. *Suppose the stage game is a decision problem. Let λ and two theories Θ_A, Θ_B be given, where Θ_A is correctly specified. Suppose there exists at least one approachable EZ with $p_A = 1$, and Θ_A is strongly identified in all such equilibria. Then Θ_A evolutionarily stable under λ -matching against Θ_B .*

The result that a resident correct specification is immune to invasions from misspecifications echoes related results in [Fudenberg and Lanzani \(2020\)](#) and [Frick, Iijima, and Ishii](#)

(2021). We focus instead on stage games where multiple agents' actions jointly determine their payoffs and characterize which misspecifications can invade a rational society in which environments.

2.7 Equilibrium Zeitgeist with Strategic Uncertainty

For most of our applications, we focus on misspecified learning about fundamental parameters of the game. But in some other environments, agents may start with a dogmatically wrong theory about others' strategies, and their equilibrium mislearning about others' behavior is the key driver of welfare.

Our setup of equilibrium zeitgeists can be adapted to allow for such strategic uncertainty. An *extended model* is a triplet (a_A, a_B, F) with $a_A, a_B \in \mathbb{A}$ and $F : \mathbb{A}^2 \rightarrow \Delta(\mathbb{Y})$. Each model contains a conjecture a_A about how group A opponents act when matched with the agent, a conjecture a_B about how group B opponents act, and a conjecture F about how strategy profiles translate into consequences for the agent. Assume each F , like F^\bullet , is given by a density or probability mass function $f(a_i, a_{-i}) : \mathbb{Y} \rightarrow \mathbb{R}_+$ for every $(a_i, a_{-i}) \in \mathbb{A}^2$. An *extended theory* $\bar{\Theta}$ is a collection of extended models: that is, a subset of $\mathbb{A}^2 \times (\Delta(\mathbb{Y}))^{\mathbb{A}^2}$. We assume the marginal of the extended theory on $(\Delta(\mathbb{Y}))^{\mathbb{A}^2}$ is metrizable. An extended theory $\bar{\Theta}$ is *correctly specified* if $\bar{\Theta} \supseteq \mathbb{A}^2 \times \{F^\bullet\}$, so the agent can make unrestricted inferences about others' play and does not rule out the correct fundamental environment F^\bullet .

Definition 6. A *zeitgeist with strategic uncertainty* $\bar{\mathfrak{Z}} = (\bar{\Theta}_A, \bar{\Theta}_B, \mu_A, \mu_B, p, \lambda, a)$ consists of: (1) two extended theories $\bar{\Theta}_A$ and $\bar{\Theta}_B$; (2) a belief over extended models for each extended theory, $\mu_A \in \Delta(\bar{\Theta}_A)$ and $\mu_B \in \Delta(\bar{\Theta}_B)$; (3) relative sizes of the two groups in the society, $p = (p_A, p_B)$ with $p_A, p_B \geq 0$, $p_A + p_B = 1$; (4) a matching assortativity parameter $\lambda \in [0, 1]$; (5) each group's strategy when matched against each other group, $a = (a_{AA}, a_{AB}, a_{BA}, a_{BB})$ where $a_{g,g'} \in \mathbb{A}$ is the strategy that an adherent of $\bar{\Theta}_g$ plays against an adherent of $\bar{\Theta}_{g'}$.

In this environment with both fundamental uncertainty and strategic uncertainty, agents still entertain some belief over their conjectures (both about others' play and about parameters of the stage game). An agent's subjective belief about the distribution of consequences in a match depends on both of these dimensions. In equilibrium, the "best-fitting" beliefs are the conjectures that lead to subjective consequence distributions closest to the true ones in the sense of KL divergence.

Definition 7. A zeitgeist with strategic uncertainty $\bar{\mathfrak{Z}} = (\bar{\Theta}_A, \bar{\Theta}_B, \mu_A, \mu_B, p, \lambda, a)$ is an *equilibrium zeitgeist with strategic uncertainty (EZ-SU)* if for every $g, g' \in \{A, B\}$, $a_{g,g'} \in \arg \max_{\hat{a} \in \mathbb{A}} \mathbb{E}_{(a_A, a_B, F) \sim \mu_g} \left[\mathbb{E}_{y \sim F(\hat{a}, a_{g'})} (\pi(y)) \right]$ and, for every $g \in \{A, B\}$, the belief μ_g is supported

on

$$\arg \min_{(\hat{a}_A, \hat{a}_B, \hat{F}) \in \bar{\Theta}_g} \left\{ \begin{array}{l} (\lambda + (1 - \lambda)p_g) \cdot D_{KL}(F^\bullet(a_{g,g}, a_{g,g}) \parallel \hat{F}(a_{g,g}, \hat{a}_g)) \\ + (1 - \lambda)(1 - p_g) \cdot D_{KL}(F^\bullet(a_{g,-g}, a_{-g,g}) \parallel \hat{F}(a_{g,-g}, \hat{a}_{-g})) \end{array} \right\}$$

where $-g$ means the group other than g .

Approachability and evolutionarily stability with strategic uncertainty are defined analogously to a world where we require agents to know each others' strategies in equilibrium.

Definition 8. An EZ-SU $\bar{\mathfrak{Z}} = (\bar{\Theta}_A, \bar{\Theta}_B, \mu_A, \mu_B, p, \lambda, a)$ with $p = (1, 0)$ is *approachable* if there exists a sequence of EZ-SUs $\bar{\mathfrak{Z}}^{(n)} = (\bar{\Theta}_A, \bar{\Theta}_B, \mu_A^{(n)}, \mu_B^{(n)}, (p_A^{(n)}, p_B^{(n)}), \lambda, (a_{AA}^{(n)}, a_{AB}^{(n)}, a_{BA}^{(n)}, a_{BB}^{(n)}))$, where $p_B^{(n)} > 0$ for all n , $p_B^{(n)} \rightarrow 0$, $\mu_A^{(n)} \rightarrow \mu_A$, $\mu_B^{(n)} \rightarrow \mu_B$, $a^{(n)} \rightarrow a$.

Definition 9. Suppose there exists at least one approachable EZ-SU with extended theories $\bar{\Theta}_A, \bar{\Theta}_B$, $p = (1, 0)$, and matching assortativity λ . Say $\bar{\Theta}_A$ is *evolutionarily stable [fragile] with strategic uncertainty* against $\bar{\Theta}_B$ under λ -matching if in all such approachable EZ-SUs, $\bar{\Theta}_A$ has a weakly higher [strictly lower] fitness than $\bar{\Theta}_B$.

3 Learning Channel and New Stability Phenomena

This section focuses on how the framework's learning channel leads to new stability phenomena.

A key feature of our theory-evolution framework is that each agent interprets her observations through the lens of her theory, thus drawing inferences about her environment (i.e., game parameters). These inferences, in turn, shape her preference over strategy profiles in the stage game. So, the learning channel endogenously determines the preferences that the adherents of different theories hold in the stage game. By contrast, the literature on preference evolution discussed in Section 1.3 precludes such inferences and endows each agent with a fixed preference.

We first show how preference evolution is embedded as a special case of our framework. We then explore the implications of the learning channel for evolutionary stability, showing that some novel stability phenomena can only arise with theory evolution, and not with preference evolution. Some of the results in our applications (e.g., Proposition 7) also show that predictions about evolutionary stability change drastically without the learning channel.

A theory Θ is called a *singleton* if $\Theta = \{F\}$ for some $F : \mathbb{A}^2 \rightarrow \Delta(\mathbb{Y})$. An agent with a singleton theory does not entertain fundamental uncertainty: she is sure that the stage game is described by F . We can view every singleton theory as a subjective utility function in the stage game, $(a_i, a_{-i}) \mapsto U_i(a_i, a_{-i}; F)$ with $U_i(a_i, a_{-i}; F) := \mathbb{E}_{y \sim F(a_i, a_{-i})}[\pi(y)]$. An EZ in a

society where all agents have singleton theories correspond to an equilibrium in a setting with preference evolution. The adherents of $\Theta_g = \{F_g\}$ hold the subjective preference $U_i(\cdot, \cdot; F_g)$ in the stage game, and all agents maximize their subjective utilities in all match types.

In a society with matching assortativity λ , an adherent of a theory with population proportion p_g is matched up with someone from the same group with probability $\lambda + (1 - \lambda)p_g$. This matching probability is an increasing and linear function in each of λ and p_g . Suppose the two subjective preferences $U_i(\cdot, \cdot; F_A)$ and $U_i(\cdot, \cdot; F_B)$ associated with the two singleton theories Θ_A and Θ_B in a society induce a unique equilibrium in matches between groups g and g' for all $g, g' \in \{A, B\}$. Then, the fitness of each theory changes linearly as we change the matching assortativity or population shares. This linearity underlies the key distinction between preference evolution and theory evolution.

Every non-singleton theory may be thought of as *a set of* preferences over stage game strategy profiles, viewing each model $F : \mathbb{A}^2 \rightarrow \Delta(\mathbb{Y})$ as one such preference. As matching assortativity or population shares change, each agent encounters a different distribution over opponent strategies. This may lead a misspecified agent to draw a different inference about the stage game parameters and may change the agent's best-response function. By contrast, in a world of preference evolution, a game between two agents with a given pair of subjective preferences always plays out in the same way, regardless of the social composition or matching assortativity of the larger society where the game takes place.

We exhibit two stability phenomena that only happen with non-singleton theories.

3.1 Stability Reversals

Stability reversal refers to a strong kind of multiplicity in the relative stability of two theories Θ_A and Θ_B under uniform matching. Recall that in an EZ, the fitness of a theory is the objective expected payoffs of its adherents, where this expectation averages across expected payoffs in matches against each of the two groups. Let a theory's *conditional fitness against group g* refer to the expected payoff of the theory's adherents in matches against group g .

Definition 10. Two theories Θ_A, Θ_B exhibit *stability reversal* if (i) in every EZ with $\lambda = 0$ and $(p_A, p_B) = (1, 0)$, Θ_A has strictly higher conditional fitness than Θ_B against group A opponents and against group B opponents, but also (ii) in every EZ with $\lambda = 0$ and $(p_A, p_B) = (0, 1)$, Θ_B has strictly higher fitness than Θ_A .

If at least one EZ is approachable with $\lambda = 0$, $(p_A, p_B) = (1, 0)$, then the first part in the definition of stability reversal is stronger than requiring Θ_A to be evolutionarily stable against Θ_B . It imposes the more stringent condition that Θ_A outperforms Θ_B not only on average, but also conditional on the opponent's group. The linearity of fitness in population

share discussed above then implies that stability reversal cannot take place if both theories are singletons (i.e., if we are in the world of preference evolution).

Proposition 2. *Two singleton theories (i.e., two subjective preferences in the stage game) cannot exhibit stability reversal in any stage game.*

Stability reversal is unique to the world of theory evolution. For an example, consider a two-player investment game where player i chooses an investment level $a_i \in \{1, 2\}$. A random productivity level P is realized according to $b^\bullet(a_i + a_{-i}) + \epsilon$ where ϵ is a zero-mean noise term, $b^\bullet > 0$. Player i gets $a_i \cdot P - 1_{\{a_i=2\}} \cdot c$. So P determines the marginal return on investment, and $c > 0$ is the cost for choosing the higher investment level, with the cost of the lower investment level normalized to 0. At the end of the game, players observe $y = (a_i, a_{-i}, P)$. The payoff matrix below displays the objective expected payoffs for different investment profiles.

	1	2
1	$2b^\bullet, 2b^\bullet$	$3b^\bullet, 6b^\bullet - c$
2	$6b^\bullet - c, 3b^\bullet$	$8b^\bullet - c, 8b^\bullet - c$

Condition 1. $5b^\bullet < c < 6b^\bullet$.

Condition 1 ensures that $a_i = 1$ is a strictly dominant strategy in the stage game, and the investment profile (2,2) Pareto dominates the investment profile (1,1). Higher investment has a positive externality as it also increases opponent's productivity.

Consider two theories in the society. Theory Θ_A is a correctly specified singleton – its adherents understand how investment profiles translate into distributions over productivity. Theory Θ_B wrongly stipulates $P = b(x_i + x_{-i}) - m + \epsilon$, where $m > 0$ is a fixed parameter of the theory and $b \in \mathbb{R}$ is a parameter that the adherents infer. We require the following condition, which is satisfied whenever $m > 0$ is large enough — that is, Θ_B is sufficiently misspecified.

Condition 2. $c < 4b^\bullet + \frac{1}{3}m$ and $c < 5b^\bullet + \frac{1}{4}m$.

We show that in contrast to the impossibility result when all theories are singletons, in this example theories Θ_A and Θ_B exhibit stability reversal.

Example 1. In the investment game, under Condition 1 and Condition 2, Θ_A and Θ_B exhibit stability reversal.

The idea is that the adherents of Θ_B overestimate the complementarity of investments, and this overestimation is more severe when they face data generated from lower investment

profiles. As a result, the match between Θ_A and Θ_B plays out in a different way depending on which theory is resident: it results in the investment profile (1, 2) when Θ_A is resident, but results in (1, 1) when Θ_B is resident.

Let $b^*(a_i, a_{-i})$ solve $\min_{b \in \mathbb{R}} D_{KL}(F^\bullet(a_i, a_{-i}) \parallel \hat{F}(a_i, a_{-i}; b, m))$, where $F^\bullet(a_i, a_{-i})$ is the objective distribution over observations under the investment profile (a_i, a_{-i}) , and $\hat{F}(a_i, a_{-i}; b, m)$ is the distribution under the same investment profile in the model where productivity is given by $P = b(x_i + x_{-i}) - m + \epsilon$. We find that $b^*(a_i, a_{-i}) = b^\bullet + \frac{m}{a_i + a_{-i}}$. That is, adherents of Θ_B end up with different beliefs about the game parameter b depending on the behavior of their typical opponents, which in turn affects how they respond to different rival investment levels. Stability reversal hinges on the fact that when Θ_A is resident and the adherents of Θ_B always meet opponents who play $a_i = 1$, they end up with a more distorted belief about the fundamental than when Θ_B is resident.

In this example, stability reversal happens because the misspecified agents hold different beliefs about a stage-game parameter depending on which theory is resident. Also, note the stage game involves non-trivial strategic interaction between the players — the complementarity in investment levels implies an agent’s best response may vary with the rival’s strategy. Both of these turn out to be necessary conditions for stability reversal in general stage games.

Definition 11. A theory Θ is *strategically independent* if for all $\mu \in \Delta(\Theta)$, $\arg \max_{a_i \in \mathbb{A}} U_i(a_i, a_{-i}; \mu)$ is the same for every $a_{-i} \in \mathbb{A}$.

The adherents of a strategically independent theory believe that while opponent’s action may affect their utility, it does not affect their best response.

Proposition 3. *In any stage game, suppose Θ_A, Θ_B exhibit stability reversal and Θ_A is the correctly specified singleton theory. Then, the beliefs that the adherents of Θ_B hold in all EZs with $p = (1, 0)$ and the beliefs they hold in all EZs with $p = (0, 1)$ form disjoint sets. Also, Θ_B is not strategically independent.*

The first claim of Proposition 3 shows that stability reversal must operate through the learning channel. So in particular, it cannot happen if the group B agents simply have a different subjective preference in the stage game. The second claim shows that stability reversal can only happen if the misspecified agents respond differently to different rival play. In particular, it cannot happen in decision problems.

3.2 Non-Monotonic Stability in Matching Assortativity

We now turn to the role of matching assortativity on the stability of theories. In the world of preference evolution, the linearity of fitness in matching assortativity discussed before

implies that if a theory Θ_A is evolutionarily stable against a theory Θ_B both under uniform matching ($\lambda = 0$) and perfectly assortative matching ($\lambda = 1$), then the same must also hold under any intermediate level of assortativity $\lambda \in (0, 1)$.

Proposition 4. *Suppose Θ_A, Θ_B are singleton theories (i.e., subjective preferences in the stage game) and Θ_A is evolutionarily stable against Θ_B with λ -matching for both $\lambda = 0$ and $\lambda = 1$. Then, Θ_A is also evolutionarily stable against Θ_B with λ -matching for any $\lambda \in [0, 1]$.*

This result does not always hold with non-singleton general theories. We use an example to show that stability need not be monotonic in matching assortativity. In this example, a correctly specified singleton theory is evolutionarily stable against another misspecified theory both when $\lambda = 0$ and when $\lambda = 1$, but it is also evolutionarily fragile for some intermediate values of λ .

Consider a stage game where each player chooses an action from $\{a_1, a_2, a_3\}$. Every player then receives a random prize, $y \in \{g, b\}$, which are worth utilities $\pi(g) = 1$, $\pi(b) = 0$. The payoff matrix below displays the objective expected utilities associated with different action profiles, which also correspond to the probabilities that the row and column players receive the good prize g .

	a_1	a_2	a_3
a_1	0.25, 0.25	0.50, 0.20	0.70, 0.15
a_2	0.20, 0.50	0.40, 0.40	0.40, 0.20
a_3	0.15, 0.70	0.20, 0.40	0.20, 0.20

Let Θ_A be the correctly specified singleton theory. The action a_1 is strictly dominant under the objective payoffs, so an adherent of Θ_A always plays a_1 in all matches. Let Θ_B be a misspecified theory $\Theta_B = \{F_H, F_L\}$. Each model F_H, F_L stipulates that the prize g is generated the probabilities in the following table, where b and c are parameters that depend on the model. The model F_H has $(b, c) = (0.8, 0.2)$ and F_L has $(b, c) = (0.1, 0.4)$.

	a_1	a_2	a_3
a_1	0.10, 0.10	0.10, c	0.10, 0.15
a_2	c , 0.10	b , b	b , 0.20
a_3	0.15, 0.10	0.20, b	0.20, 0.20

The learning channel for the biased mutants leads the correctly specified theory to have non-monotonic evolutionarily stability in terms of matching assortativity.

Example 2. In this stage game, Θ_A is evolutionarily stable against Θ_B under λ -matching when $\lambda = 0$ and $\lambda = 1$, but it is also evolutionarily fragile under λ -matching when $\lambda \in (\lambda_l, \lambda_h)$, where $0 < \lambda_l < \lambda_h < 1$ are $\lambda_l = 0.25$, $\lambda_h \approx 0.56$.

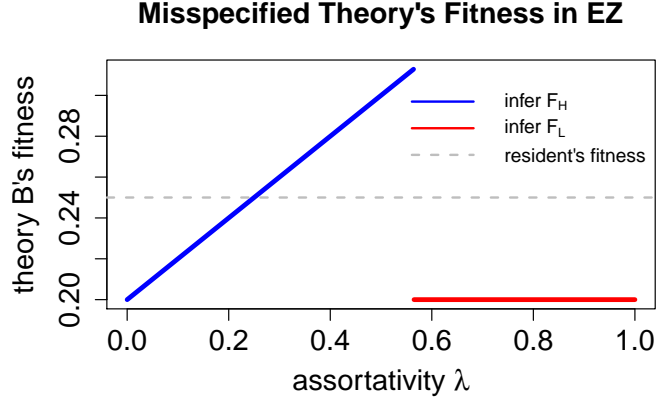


Figure 1: The EZ fitness of Θ_B for different values of matching assortativity λ when $p_B = 0$. (The EZ fitness of the resident theory Θ_A is always 0.25.) In the blue region, there is a unique EZ where the adherents of Θ_B infer F_H and receive linearly increasing average payoffs across all matches as λ increases. In the red region, there is an EZ where the adherents of Θ_B infer F_L and receive payoff 0.2 in all matches, regardless of λ .

To understand the intuition, examine the match between two adherents of Θ_B . If they believe in F_H , they will play the action profile (a_2, a_2) and generate the objective payoff profile $(0.4, 0.4)$, a Pareto improvement compared to the correctly specified outcome (a_1, a_1) . The problem is that the data generated from the (a_2, a_2) profile provides a better fit for F_L than F_H , since the objective 40% probability of getting prize g is closer to F_L 's conjecture of 10% than F_H 's conjecture of 80%. A belief in F_H — and hence the profile (a_2, a_2) — cannot be sustained if the mutants only play each other. On the other hand, when an adherent of Θ_B plays a correctly specified Θ_A adherent, both models F_H and F_L prescribe a best response of a_2 against the Θ_A adherent's play a_1 . The data generated from the (a_2, a_1) profile lead biased agents to the model F_H that enables cooperative behavior within the mutant community. But, these matches against correctly specified opponents harm the mutant's welfare, as they only get an objective payoff of 0.2.

Therefore, the most advantageous interaction structure for the mutants is one where they can calibrate the model F_H using the data from matches against correctly specified opponents, then extrapolate this optimistic belief about b to coordinate on (a_2, a_2) in matches against fellow mutants. This requires the mutants to match with intermediate assortativity. Figure 1 depicts the equilibrium fitness of the mutant theory Θ_B as a function of assortativity. While payoffs of Θ_B adherents increase in λ at first, eventually they drop when mutant-vs-mutant matches become sufficiently frequent that a belief in F_H can no longer be sustained. The preference evolution framework does not allow this non-linear and even non-monotonic change in fitness with respect to λ , which the theory evolution framework accommodates.

While this analysis fixes the population sizes at $p_A = 1$, $p_B = 0$ and varies the matching assortativity, we find a similar conclusion if we fix the matching assortativity and vary the population sizes. For instance, let assortativity be $\lambda = 0.5$ and consider the equilibrium fitness of the misspecified theory Θ_B as a function of its population share $p_B \in [0, 1]$. The probability that a misspecified agent meets another misspecified agent is $0.5 + 0.5p_B$, which is the same as when the population sizes are $p_A = 1$, $p_B = 0$ and $\lambda = 0.5 + 0.5p_B$. By this equivalence, it is easy to show that the fitness of Θ_B is non-monotonic in p_B . In particular, Θ_B has strictly higher EZ fitness than Θ_A whenever $p_B \leq 0.128$, but strictly lower EZ fitness than Θ_A when p_B is high enough.

4 Higher-Order Misspecifications in Linear-Quadratic-Normal Games

We apply our framework to study the stability of misperceptions of the information structure in linear quadratic normal (LQN) games. LQN games have been used as a tractable workhorse model for studying comparative statics of equilibrium outcomes with respect to changes in information (e.g., [Bergemann and Morris \(2013\)](#)). In this application, we exploit the same tractability to study the evolutionary stability of correct beliefs about the information structure to misspecifications — in particular, misspecifications about the correlation in information between different players. The key conclusion is that a society of rational residents with correct beliefs about how private signals are correlated is evolutionarily fragile against misspecified mutants who suffer from either correlation neglect or projection bias. The type of bias that gets selected depends on the matching assortativity λ in the society.

4.1 Stage Game and Misperceptions of Information Structure

In the LQN setup we consider, we interpret the players as competing firms that possess correlated private information about market demand. At the start of the stage game, Nature's type (i.e., a demand state) ω is drawn from $\mathcal{N}(0, \sigma_\omega^2)$, where $\mathcal{N}(\mu, \sigma^2)$ is the normal distribution with mean μ and variance σ^2 . Each of the two players i (i.e., firms) receives a private signal $s_i = \omega + \epsilon_i$, then chooses an action $q_i \in \mathbb{R}$ (i.e., a quantity). Market price is then realized according to $P = \omega - r^\bullet \cdot \frac{1}{2}(q_1 + q_2) + \zeta$, where $\zeta \sim \mathcal{N}(0, (\sigma_\zeta^\bullet)^2)$ is an idiosyncratic price shock that is independent of all the other random variables. Firm i 's profit in the game is $q_i P - \frac{1}{2}q_i^2$.

The stage game is parametrized by the strictly positive terms σ_ω^2 , r^\bullet , and $(\sigma_\zeta^\bullet)^2$, which represent variance in market demand, the elasticity of market price with respect to average

quantity supplied, and the variance of price shocks. These parameters remain constant through all matches. But in every match, demand state ω , signals (s_i) , and price shock ζ are redrawn, independently across matches. The environment can be interpreted as a market with daily fluctuations in demand, but the fluctuations are generated according to a fixed set of fundamental parameters.

In the LQN game, market prices and quantity choices may be positive or negative. To interpret, when $P > 0$, the market pays for each unit of good supplied, and market price decreases in total supply. When $P < 0$, the market pays for disposal of the good. Firms make money by submitting negative quantities, which represent offers to remove the good from the market. The per-unit disposal fee decreases as the firms offer to dispose more. The cost $\frac{1}{2}q_i^2$ represents either a convex production cost or a convex disposal cost, depending on the sign of q_i .

We now turn to the information structure of the stage game — that is, the joint distribution of (ω, s_i, s_{-i}) . The firms' signals $s_i = \omega + \epsilon_i$ are conditionally correlated given ω . The error terms ϵ_i are generated by

$$\epsilon_i = \frac{\kappa}{\sqrt{\kappa^2 + (1 - \kappa)^2}}z + \frac{1 - \kappa}{\sqrt{\kappa^2 + (1 - \kappa)^2}}\eta_i,$$

where $\eta_i \sim \mathcal{N}(0, \sigma_\epsilon^2)$ is the idiosyncratic component of the error generated i.i.d. across i , and $z \sim \mathcal{N}(0, \sigma_\epsilon^2)$ is the common component for both i . Here, $\kappa \in [0, 1]$ parametrizes the conditional correlation of the two firms' signals. Higher κ leads to an information structure with higher conditional correlation. When $\kappa = 0$, s_i and s_{-i} are conditionally uncorrelated given the state (though still unconditionally correlated since both depend on ω). When $\kappa = 1$, we always have $s_i = s_{-i}$. The functional form of ϵ_i ensures the variance of the signals $\text{Var}(s_i)$ remains constant across all possible values of κ .

We consider a family of misspecifications about the information structure parametrized by misperceptions of κ . The objective information structure is given by $\kappa = \kappa^\bullet$. Note that a misspecified information structure associated with a wrong κ leads to a higher-order misspecification about the state ω in the stage game. Suppose agents are correct about the distributions of ω , η_i , and z . Write \mathbb{E}_κ for expectation under the information structure with correlational parameter κ . Then $\mathbb{E}_\kappa[\omega \mid s_i]$ is the same for all κ — in particular, even an agent who believes in some $\kappa \neq \kappa^\bullet$ makes a correct first-order inference about the expectation of the market demand, given her own information. But, one can show (Lemma 1) there exists a strictly increasing and strictly positive function $\psi(\kappa)$ so that $\mathbb{E}_\kappa[s_{-i} \mid s_i] = \psi(\kappa) \cdot s_i$ for all $s_i \in \mathbb{R}, \kappa \in [0, 1]$. The misspecified agent holds a wrong belief about the rival's signal, and thus a wrong belief about the rival's belief about ω .

Many experiments have found that subjects do not form accurate beliefs about the beliefs of others. We draw a connection between the misperception we study and the statistical biases that have been previously documented:

Definition 12. Let $\tilde{\kappa}$ be a player’s perceived κ . A player suffers from *correlation neglect* if $\tilde{\kappa} < \kappa^\bullet$. A player suffers from *projection bias* if $\tilde{\kappa} > \kappa^\bullet$.

Under correlation neglect, agents believe signals are more independent from one another than they really are. Under projection bias, agents “project” their own information onto others and exaggerate the similarity between others’ signals and their own signals. We are agnostic about the origin of these misspecifications about correlation. They may be psychological in nature and come directly from the agents’ cognitive biases, or they could be driven by more complex mechanisms.⁵ We instead ask whether such misspecifications could persist in the society once they appear.

4.2 Formalizing Strategies and Theories

We translate the environment described above into the formalism from Section 2.

A strategy in the stage game is a function $Q_i : \mathbb{R} \rightarrow \mathbb{R}$ that assigns a quantity $Q_i(s_i)$ to every signal s_i . The strategy is called *linear* if there exists an $\alpha_i \geq 0$ so that $Q_i(s_i) = \alpha_i s_i$ for every $s_i \in \mathbb{R}$. We will later show that the best response to any linear strategy is linear, regardless of the agent’s belief about the correlation parameter and market price elasticity (Lemma 2). We therefore restrict attention to linear strategies and let $\mathbb{A} = [0, \bar{M}_\alpha]$ for some $\bar{M}_\alpha < \infty$, where a typical element $\alpha_i \in \mathbb{A}$ corresponds to the linear strategy with coefficient α_i .

We suppose all parameters of the stage game are common knowledge except for r^\bullet , κ^\bullet , and σ_ζ^\bullet . To investigate the evolutionary implications of higher-order misspecifications about the state, we consider theories that are dogmatic and possibly wrong about κ , but allow agents to make inferences about r and σ_ζ . We let the space of consequences be $\mathbb{Y} = \mathbb{R}^3$, where a typical consequence $y = (s_i, q_i, P)$ shows the agent’s signal, quantity choice, and the market price. The consequence y delivers the utility $\pi(y) := q_i P - \frac{1}{2} q_i^2$. We consider theories parametrized by κ , $\Theta(\kappa) := \{F_{r,\kappa,\sigma_\zeta} : r \in [0, \bar{M}_r], \sigma_\zeta \in [0, \bar{M}_{\sigma_\zeta}]\}$ for some $\bar{M}_r, \bar{M}_{\sigma_\zeta} < \infty$. So each $\Theta(\kappa)$ is a set of conjectures of the game environment indexed by the parameters $(r, \kappa, \sigma_\zeta)$, but all reflecting a dogmatic belief in the correlation parameter κ . Each $F_{r,\kappa,\sigma_\zeta} : \mathbb{A} \times \mathbb{A} \rightarrow \Delta(\mathbb{Y})$ is

⁵For example, Hansen, Misra, and Pai (2021) show that multiple agents simultaneously conducting algorithmic price experiments in the same market may generate correlated information which get misinterpreted as independent information, a form of correlation neglect for firms. Goldfarb and Xiao (2019) structurally estimate a model of thinking cost and find that bar owners over-extrapolate the effect of today’s weather shock on future profitability.

such that $F_{r,\kappa,\sigma_\zeta}(\alpha_i, \alpha_{-i})$ gives the distribution over i 's consequences in a stage game with parameters $(r, \kappa, \sigma_\zeta)$, when i uses the linear strategy α_i against an opponent using the linear strategy α_{-i} .

While agents learn about both r and σ_ζ , it is their (mis)inferences about the market price elasticity r that drives the main results. Since each firm's profit is linear in the market price, an agent's belief about the variance of the idiosyncratic price shock does not change her expected payoffs or behavior. We use inference over σ_ζ to simplify our analysis: this parameter absorbs changes in the variance of market price under different correlation structures. A Bayesian agent whose data are all generated from the same strategy profile only learn about r using the mean of the market price in the data, not its variance.

In formalizing the stage game and translating misperceptions of the information structure into theories, we have assumed that the space of feasible linear strategies $\alpha_i \in [0, \bar{M}_\alpha]$ and the domain of inference over game parameters $r \in [0, \bar{M}_r], \sigma_\zeta \in [0, \bar{M}_{\sigma_\zeta}]$ are bounded sets. These compactness assumptions help ensure that EZ exist. In analyzing evolutionary stability, we will focus on the case where the bounds $\bar{M}_\alpha, \bar{M}_r, \bar{M}_{\sigma_\zeta}$ are finite but sufficiently large, so that the optimal behavior and beliefs are interior. We introduce the following shorthand:

Notation 1. A result is said to hold “*with high enough price volatility and large enough strategy space and inference space*” if, whenever the strategy space $[0, \bar{M}_\alpha]$ has $\bar{M}_\alpha \geq \frac{1/\sigma_\epsilon^2}{1/\sigma_\epsilon^2 + 1/\sigma_\omega^2}$, there exist $0 < L_1, L_2, L_3 < \infty$ so that for any objective game F^\bullet with $(\sigma_\zeta^\bullet)^2 \geq L_1$ and with theories where the parameter spaces $r \in [0, \bar{M}_r], \sigma_\zeta \in [0, \bar{M}_{\sigma_\zeta}]$ are such that $\bar{M}_{\sigma_\zeta}^2 \geq (\sigma_\zeta^\bullet)^2 + L_2$ and $\bar{M}_r \geq L_3$, the result is true.

4.3 Subjective Best Response and Misspecified Inference

In order to determine which theories (i.e., perceptions of κ) are stable against rival theories, we must characterize the relevant equilibrium zeitgeists. This section develops a number of preliminary results that relate beliefs about the game parameters to best responses, and conversely strategy profiles to the KL-divergence minimizing inferences.

We begin by proving the result alluded to earlier: under normality, every agent's inferences about the state and about opponent's signal are linear functions of her own signal. The linear coefficient on the latter increases with the correlation parameter κ .

Lemma 1. *There exists a strictly increasing function $\psi(\kappa)$, with $\psi(0) > 0$ and $\psi(1) = 1$, so that $\mathbb{E}_\kappa[s_{-i} | s_i] = \psi(\kappa) \cdot s_i$ for all $s_i \in \mathbb{R}, \kappa \in [0, 1]$. Also, there exists a strictly positive $\gamma \in \mathbb{R}$ so that $\mathbb{E}_\kappa[\omega | s_i] = \gamma \cdot s_i$ for all $s_i \in \mathbb{R}, \kappa \in [0, 1]$.*

Linearity of $\mathbb{E}[\omega | s_i]$ and $\mathbb{E}[s_{-i} | s_i]$ in s_i allows us explicitly characterize the corresponding linear best responses, given beliefs about κ and elasticity r . For Q_i, Q_{-i} (not necessarily

linear) strategies in the stage game and $\mu \in \Delta(\Theta(\kappa))$, let $U_i(Q_i, Q_{-i}; \mu)$ be i 's subjective expected utility from playing Q_i against Q_{-i} , under the belief μ .

Lemma 2. *For α_{-i} a linear strategy,*

$$U_i(\alpha_i, \alpha_{-i}; \mu) = \mathbb{E}[s_i^2] \cdot \left(\alpha_i \gamma - \frac{1}{2} \hat{r} \alpha_i^2 - \frac{1}{2} \hat{r} \psi(\kappa) \alpha_i \alpha_{-i} - \frac{1}{2} \alpha_i^2 \right)$$

for every linear strategy α_i , where $\hat{r} = \int r \, d\mu(r, \kappa, \sigma_\zeta)$ is the mean of μ 's marginal on elasticity. For $\kappa \in [0, 1]$ and $r > 0$, $\alpha_i^{BR}(\alpha_{-i}; \kappa, r) := \frac{\gamma - \frac{1}{2} r \psi(\kappa) \alpha_{-i}}{1+r}$ best responds to α_{-i} among all strategies $Q_i : \mathbb{R} \rightarrow \mathbb{R}$ for all $\sigma_\zeta > 0$.

Lemma 2 shows that $\alpha_i^{BR}(\alpha_{-i}; \kappa, r)$ is not only the best-responding linear strategy when opponent plays α_{-i} and i believes in correlation parameter κ and elasticity r , it is also optimal among the class of all strategies $Q_i(s_i)$ against the same opponent play and under the same beliefs.

Call a linear strategy more *aggressive* if its coefficient $\alpha_i \geq 0$ is larger. One implication of Lemma 2 is that agent i 's subjective best response function becomes more aggressive when i believes in lower κ or lower r . We have $\frac{\partial \alpha_i^{BR}}{\partial \kappa} < 0$ because the agent can better capitalize on her private information about market demand when her rival does not share the same information. We have $\frac{\partial \alpha_i^{BR}}{\partial r} < 0$ because the agent can be more aggressive in general when facing an inelastic market price.

We now turn to equilibrium inference about the market price elasticity r^\bullet . The following lemma shows that any linear profile generates data whose KL-divergence can be minimized to 0 by a unique value of r . We also characterize how this inference about elasticity depends on the strategy profiles and the agent's belief about the correlation parameter κ . As mentioned earlier, we focus on the case where the bounds on the inferences $r \in [0, \bar{M}_r]$, $\sigma_\zeta \in [0, \bar{M}_{\sigma_\zeta}]$ are sufficiently large to ensure that the KL-divergence minimization problem is well-behaved.

Lemma 3. *With high enough price volatility and large enough strategy space and inference space, for every $\alpha_i, \alpha_{-i} \in [0, \bar{M}_\alpha]$, we have $D_{KL}(F_{r^\bullet, \kappa^\bullet, \sigma_\zeta^\bullet}(\alpha_i, \alpha_{-i}) \parallel F_{\hat{r}, \kappa, \hat{\sigma}_\zeta}(\alpha_i, \alpha_{-i})) = 0$ for exactly one pair $\hat{r} \in [0, \bar{M}_r]$, $\hat{\sigma}_\zeta \in [0, \bar{M}_{\sigma_\zeta}]$. This \hat{r} is given by $r_i^{INF}(\alpha_i, \alpha_{-i}; \kappa^\bullet, \kappa, r^\bullet) := \frac{r^\bullet \alpha_i + \alpha_{-i} \psi(\kappa^\bullet)}{\alpha_i + \alpha_{-i} \psi(\kappa)}$.*

Lemma 3 implies that an agent's inference about r is strictly decreasing in her belief about the correlation parameter κ . To understand why, assume player i uses the linear strategy α_i and player $-i$ uses the linear strategy α_{-i} . After receiving a private signal s_i , player i expects to face a price distribution with a mean of $\gamma s_i - r(\frac{1}{2} \alpha_i s_i + \frac{1}{2} \alpha_{-i} \mathbb{E}_\kappa[s_{-i} | s_i])$. Under projection bias $\kappa > \kappa^\bullet$, $\mathbb{E}_\kappa[s_{-i} | s_i]$ is excessively steep in s_i . For example, following a large and positive s_i , the agent overestimates the similarity of $-i$'s signal and wrongly predicts

that $-i$ must also choose a very high quantity, and thus becomes surprised when market price remains high. The agent then wrongly infers that the market price elasticity must be low. Therefore, in order to rationalize the average market price conditional on own signal, an agent with projection bias must infer $r < r^\bullet$. For similar reasons, an agent with correlation neglect infers $r > r^\bullet$.

Combining Lemma 2 and Lemma 3, we find that increasing κ has an *a priori* ambiguous impact on the agent's equilibrium aggressiveness. Increasing κ has the direct effect of lowering aggression (by Lemma 2), but it also causes the indirect effect of lowering inference about r (by Lemma 3) and therefore increases aggression (by Lemma 2). Nevertheless, we show in the results below that the indirect effect through the mislearning channel dominates, and the evolutionary stability of correlational errors are driven by this channel. We show in Section 4.6 that the results are reversed when we shut down the learning channel.

Lemma 3 considers the problem of KL-divergence minimization when all of the data are generated from a single strategy profile, $(\alpha_{-i}, \alpha_{-i})$. It implies that if $\lambda \in \{0, 1\}$ and $(p_A, p_B) = (1, 0)$, that is matching is either perfectly uniform or perfectly assortative in a homogeneous society, then every agent can find a model to exactly fit her equilibrium data. This is because agents only match with opponents from one group in the EZ. The self-confirming property lends a great deal of tractability and allows us to provide sharp comparative statics and assess the stability of theories.

With interior population shares, agents can observe consequences from matches against the adherents of both Θ_A and Θ_B . Thus, they must find a single set of parameters for the stage game that best fits all of their data, and even this best-fitting model will have positive KL divergence in equilibrium. The next lemma shows the LQN game satisfies the sufficient conditions from Appendix B (Assumptions A.1 through A.5) for the existence and upper hemicontinuity of EZs. So, the tractable analysis in homogeneous societies remains robust to the introduction of a small but non-zero share of a mutant theory.

Lemma 4. *For every $r^\bullet, \sigma_\zeta^\bullet \geq 0$, $\lambda \in [0, 1]$, $\kappa^\bullet, \kappa \in [0, 1]$, $\bar{M}_\alpha, \bar{M}_{\sigma_\zeta}, \bar{M}_r < \infty$, the LQN with objective parameters $(r^\bullet, \kappa^\bullet, \sigma_\zeta^\bullet)$, strategy space $\mathbb{A} = [0, \bar{M}_\alpha]$ and theories $\Theta(\kappa^\bullet), \Theta(\kappa)$ with parameter spaces $[0, \bar{M}_r], [0, \bar{M}_{\sigma_\zeta}]$ satisfy Assumptions A.1, A.2, A.3, A.4, and A.5. Therefore, EZs in LQN are upper hemicontinuous in population sizes.*

4.4 Uniform Matching ($\lambda = 0$) and Projection Bias

We now describe our main results on the evolutionary instability of correctly specified beliefs about the information structure. Our first main result is that in a society where agents are uniformly matched, a correctly specified κ will be evolutionarily fragile against some amount of projection bias. The proof of this result involves characterizing the *asymmetric*

equilibrium strategy profile in matches between the correctly specified residents and the projection-biased mutants, and proving that a small amount of projection bias leads the mutants to have higher payoffs in the resident-vs-mutant matches than the residents' payoffs in the resident-vs-resident matches.

Proposition 5. *Let $r^\bullet > 0$, $\kappa^\bullet \in [0, 1]$ be given. With high enough price volatility and large enough strategy space and inference space, there exist $\underline{\kappa} < \kappa^\bullet < \bar{\kappa}$ so that in societies with two theories $(\Theta_A, \Theta_B) = (\Theta(\kappa^\bullet), \Theta(\kappa))$ where $\kappa \in [\underline{\kappa}, \bar{\kappa}]$, there is a unique EZ with uniform matching ($\lambda = 0$) and $(p_A, p_B) = (1, 0)$. The equilibrium fitness of $\Theta(\kappa)$ is strictly higher than that of $\Theta(\kappa^\bullet)$ if $\kappa > \kappa^\bullet$, and strictly lower if $\kappa < \kappa^\bullet$.*

Combining this result with Lemma 4, we conclude that in societies with theories $\Theta(\kappa^\bullet)$ and $\Theta(\kappa)$ where κ is slightly above κ^\bullet , the unique EZ is approachable. Hence, the correct specification is not evolutionarily stable against a small amount of projection bias.

Intuitively, as discussed after Lemma 3, projection bias generates a commitment to aggression as it leads the biased agents to under-infer market price elasticity. It is well known that in Cournot oligopoly games, such commitment can be beneficial. For instance, if quantities are chosen sequentially, the first mover obtains a higher payoff compared to the case where quantities are chosen simultaneously. A similar force is at work here, but the source of the commitment is different. Misspecification about signal correlation leads to misinference about r^\bullet , which causes the mutants to credibly respond to their opponents' play in an overly aggressive manner. The rational residents, who can identify the mutants in the population, back down and yield a larger share of the surplus. While projection bias is beneficial in small amounts, it is also intuitive that excessive aggression would be detrimental as well, as overproduction can be individually suboptimal.

4.5 Fully Assortative Matching ($\lambda = 1$) and Correlation Neglect

Turning to the case of perfectly assortative matching, we obtain the opposite result: evolutionary stability now selects for theories with correlation neglect. The fragility of the correct specification is even starker here, as we show that any level of correlation neglect leads to higher equilibrium fitness.

Proposition 6. *Let $r^\bullet > 0$, $\kappa^\bullet \in [0, 1]$ be given. With high enough price volatility and large enough strategy space and inference space, in societies with two theories $(\Theta_A, \Theta_B) = (\Theta(\kappa_A), \Theta(\kappa_B))$ where $\kappa_A \leq \kappa_B$, the fitness of Θ_A is weakly higher than that of Θ_B in every EZ with any population proportion p and perfectly assortative matching ($\lambda = 1$).*

Combining this result with Lemma 4, we conclude that under Proposition 6's conditions with $(p_A, p_B) = (1, 0)$, at least one EZ is approachable, and each theory's fitness is invariant

across all approachable EZs. Furthermore, this fitness is strictly decreasing in κ . Hence, for any $\kappa_A < \kappa_B$, theory $\Theta(\kappa_A)$ is evolutionarily stable against theory $\Theta(\kappa_B)$. Specializing to $\kappa_B = \kappa^\bullet$, we conclude that the correct specification is evolutionarily fragile against any level of correlation neglect.

As discussed after Lemma 3, correlation neglect makes agents over-infer market price elasticity, and thus lets them commit to more cooperative behavior (i.e., linear strategies with a smaller coefficient α_i). Rational opponents would take advantage of such agents, but the biased agents never match up against rational opponents in a society with perfectly assortative matching. Note also that in the uniform matching case, projection bias leads to higher payoff for the mutant at the expense of the rational opponent's payoff. With perfectly assortative matching, correlation neglect Pareto improves both biased agents' payoffs.

To understand why equilibrium fitness is a monotonically decreasing function of κ with perfectly assortative matching, let α^{TEAM} denote the symmetric linear strategy profile that maximizes the sum of the two firms' expected objective payoffs. We can show that among symmetric strategy profiles, players' payoffs strictly decrease in their aggressiveness in the region $\alpha > \alpha^{TEAM}$. We can also show that with $\lambda = 1$ and any $\kappa \in [0, 1]$, the equilibrium play among two adherents of $\Theta(\kappa)$ strictly increases in aggression as κ grows, and it is always strictly more aggressive than α^{TEAM} . Lowering perception of κ confers an evolutionary advantage by bringing play monotonically closer to the team solution α^{TEAM} in equilibrium.

4.6 The Necessity of the (Mis)Learning Channel

The key mechanism behind Proposition 5 and Proposition 6 is that misperceptions about κ confer an evolutionary advantage through the learning channel: they cause the misspecified agents to misinfer some other parameter of the stage game. This mislearning is strategically beneficial as it commits the agents to certain behavior that increases their equilibrium payoffs against their typical opponents, given the matching assortativity. Section 3 showed that the learning channel unique to the world of theory evolution permits novel stability phenomena in general games, and here we find the same channel is also indispensable for the predictions in this particular application. The results about the evolutionary fragility of the correct specification in Proposition 5 and Proposition 6 would be reversed without it.

Proposition 7. *Let $r^\bullet > 0$, $\kappa^\bullet \in [0, 1]$ be given. With high enough price volatility and large enough strategy space and inference space, there exists $\epsilon > 0$ so that for any $\kappa_l, \kappa_h \in [0, 1]$, $\kappa_l < \kappa^\bullet < \kappa_h \leq \kappa^\bullet + \epsilon$, the correctly specified theory $\Theta(\kappa^\bullet)$ is evolutionarily stable against the singleton theory $\{F_{r^\bullet, \kappa_h, \sigma_c^\bullet}\}$ under uniform matching ($\lambda = 0$), and evolutionarily stable against the singleton theory $\{F_{r^\bullet, \kappa_l, \sigma_c^\bullet}\}$ under perfectly assortative matching ($\lambda = 1$).*

In this proposition, we consider agents with singleton theories who misperceive the signal correlation structure but hold dogmatic and correct beliefs about the other game parameters, including the elasticity of market price. Once the mislearning channel is shut down, we find that misperceptions about κ that used to confer an evolutionary advantage under a certain matching assortativity can no longer invade a society of correctly specified residents.

4.7 Evolutionary Stability in Incomplete-Information Games

We turn to general incomplete-information games and provide a condition for a theory to be evolutionarily fragile against a “nearby” misspecified theory. This condition shows how assortativity and the learning channel shape the evolutionary selection of theories for a broader class of stage games and biases. We also relate the condition to the specific results studied so far in this application.

Consider a stage game where a state of the world ω is realized at the start of the game. Players 1 and 2 observe private signals $s_1, s_2 \in S \subseteq \mathbb{R}$, possibly correlated given ω . The objective distribution of (ω, s_1, s_2) is \mathbb{P}^\bullet . Based on their signals, players choose actions $q_1, q_2 \in \mathbb{R}$ and receive random consequences $y_1, y_2 \in \mathbb{Y}$. The distribution over consequences as a function of $(\omega, s_1, s_2, q_1, q_2)$ and the utility over consequences $\pi : \mathbb{Y} \rightarrow \mathbb{R}$ are such that each player i 's objective expected utility from taking action q_i against opponent action q_{-i} in state ω is given by $u_i^\bullet(q_i, q_{-i}; \omega)$, differentiable in its first two arguments.

For an interval of real numbers $[\underline{\kappa}, \bar{\kappa}]$ with $\underline{\kappa} < \bar{\kappa}$ and $\kappa^\bullet \in (\underline{\kappa}, \bar{\kappa})$, suppose there is a family of theories $(\Theta(\kappa))_{\kappa \in [\underline{\kappa}, \bar{\kappa}]}$. Fix $\lambda \in [0, 1]$ and a strategy space $\mathbb{A} \subseteq \mathbb{R}^S$, representing the feasible signal-contingent strategies. Suppose the two theories in the society are $\Theta_A = \Theta(\kappa^\bullet)$ and $\Theta_B = \Theta(\kappa)$ for some $\kappa \in [\underline{\kappa}, \bar{\kappa}]$. The next assumption requires there to be a unique EZ with $(p_A, p_B) = (1, 0)$ in such societies with any $\kappa \in [\underline{\kappa}, \bar{\kappa}]$, and further requires the EZ to feature linear equilibria. Linear equilibria exist and are unique in a large class of games outside of the duopoly framework, and in particular in LQN games under some conditions on the payoff functions (see, e.g., Angeletos and Pavan (2007)).

Assumption 1. *Suppose there is a unique EZ under λ -matching and population proportions $(p_A, p_B) = (1, 0)$ with $\Theta_A = \Theta(\kappa^\bullet)$, $\Theta_B = \Theta(\kappa)$ for every $\kappa \in [\underline{\kappa}, \bar{\kappa}]$. Suppose the κ -indexed EZ strategy profiles $(\sigma(\kappa)) = (\sigma_{AA}(\kappa), \sigma_{AB}(\kappa), \sigma_{BA}(\kappa), \sigma_{BB}(\kappa))$ are linear, i.e., $\sigma_{gg'}(\kappa)(s_i) = \alpha_{gg'}(\kappa) \cdot s_i$ with $\alpha_{gg'}(\kappa)$ differentiable in κ . Suppose that in the EZ with $\kappa = \kappa^\bullet$, $\alpha_{AA}(\kappa^\bullet)$ is objectively interim-optimal against itself.⁶ Finally, assume for every κ , Assumptions A.1, A.2, A.3, A.4, and A.5 are satisfied.*

⁶More precisely, for every $s_i \in S$, $\alpha_{AA}(\kappa^\bullet) \cdot s_i$ maximizes the agent's objective expected utility across all of \mathbb{R} when $-i$ uses the same linear strategy $\alpha_{AA}(\kappa^\bullet)$.

Proposition 8. Let $\alpha^\bullet := \alpha_{AA}(\kappa^\bullet)$. Then, under Assumption 1, if

$$\mathbb{E}^\bullet \left[\mathbb{E}^\bullet \left[\frac{\partial u_1^\bullet}{\partial q_2}(\alpha^\bullet s_1, \alpha^\bullet s_2, \omega) \cdot [(1 - \lambda)\alpha'_{AB}(\kappa^\bullet) + \lambda\alpha'_{BB}(\kappa^\bullet)] \cdot s_2 \mid s_1 \right] \right] > 0,$$

then there exists some $\epsilon > 0$ so that $\Theta(\kappa^\bullet)$ is evolutionarily fragile against theories $\Theta(\kappa)$ with $\kappa \in (\kappa^\bullet, \kappa^\bullet + \epsilon] \cap [\underline{\kappa}, \bar{\kappa}]$. Also, if

$$\mathbb{E}^\bullet \left[\mathbb{E}^\bullet \left[\frac{\partial u_1^\bullet}{\partial q_2}(\alpha^\bullet s_1, \alpha^\bullet s_2, \omega) \cdot [(1 - \lambda)\alpha'_{AB}(\kappa^\bullet) + \lambda\alpha'_{BB}(\kappa^\bullet)] \cdot s_2 \mid s_1 \right] \right] < 0,$$

then there exists some $\epsilon > 0$ so that $\Theta(\kappa^\bullet)$ is evolutionarily fragile against theories $\Theta(\kappa)$ with $\kappa \in [\kappa^\bullet - \epsilon, \kappa^\bullet) \cap [\underline{\kappa}, \bar{\kappa}]$. Here \mathbb{E}^\bullet is the expectation with respect to the objective distribution of (ω, s_1, s_2) under \mathbb{P}^\bullet .

Proposition 8 describes a general condition to determine whether a correctly specified theory is evolutionarily fragile against a nearby misspecified mutant theory. The condition asks if a slight change in the mutant theory's κ leads mutants' opponents to change their equilibrium actions such that the mutants become better off on average. These opponents are the residents under uniform matching $\lambda = 0$, so $\alpha'_{AB}(\kappa^\bullet)$ is relevant. These opponents are other mutants under perfectly assortative matching $\lambda = 1$, so $\alpha'_{BB}(\kappa^\bullet)$ is relevant.

Proposition 8 implies that one should only expect the correctly specified theory to be stable against all nearby theories in “special” cases — that is, when the expectation in the statement of Proposition 8 is exactly equal to 0. One such special case is when the agents face a decision problem where 2's action does not affect 1's payoffs, that is $\frac{\partial u_1^\bullet}{\partial q_2} = 0$. This sets the expectation to zero, so the result never implies that the correctly specified theory is evolutionarily fragile against a misspecified theory in such decision problems.

In the duopoly game analyzed previously, we have $\frac{\partial u_1^\bullet}{\partial q_2}(q_1, q_2, \omega) = -\frac{1}{2}r^\bullet q_1$. Player 1 is harmed by player 2 producing more if $q_1 > 0$, and helped if $q_1 < 0$. From straightforward algebra, the expectation in Proposition 8 simplifies to

$$\mathbb{E}^\bullet[s_1^2] \cdot \left(-\frac{1}{2}\psi(\kappa^\bullet)r^\bullet\alpha^\bullet\right) \cdot [(1 - \lambda)\alpha'_{AB}(\kappa^\bullet) + \lambda\alpha'_{BB}(\kappa^\bullet)].$$

The proof of Proposition 5 shows that when $\lambda = 0$, $\alpha'_{AB}(\kappa^\bullet) < 0$. The proof of Proposition 6 shows that when $\lambda = 1$, $\alpha'_{BB}(\kappa^\bullet) > 0$. The uniqueness of EZ also follow from these propositions, for an open interval of κ containing κ^\bullet . We restrict \mathbb{A} to the set of linear strategies, and Lemma 2 implies the linear strategies played by two correctly specified firms against each other are interim optimal. Finally, Lemma 4 verifies that Assumptions A.1 through A.5 are satisfied. Therefore, the conditions of Proposition 8 are satisfied for $\lambda \in \{0, 1\}$, and we deduce the correctly specified theory is evolutionarily fragile against slightly

higher κ (for $\lambda = 0$) and slightly lower κ (for $\lambda = 1$).

5 Evolutionary Stability of Analogy Classes

In the next application, we use the framework of theory evolution to provide a justification for coarse analogy classes in games. Jehiel (2005) introduces the solution concept of analogy-based expectation equilibrium (ABEE) in extensive-form games, where agents group opponents' nodes in an extensive-form game into *analogy classes* and only keep track of aggregate statistics of opponents' average behavior within each analogy class. An ABEE is a strategy profile where agents best respond to the belief that at all nodes in every analogy class, opponents behave according to the average behavior in the analogy class. In the ensuing literature that applies ABEE to different settings, analogy classes are usually exogenously given and interpreted as arising from coarse feedback or agents' cognitive limitations.⁷ We show through an example that we can encode analogy classes using suitably defined extended theories (that is, sets of conjectures about how others play and about parameters of the game) whose marginals on opponents' play are restricted subsets of extensive-form strategies, and that the matches between any two groups in an EZ-SU constitute ABEEs. We then investigate which analogy classes are more likely to arise by studying the stability of different theories (i.e., analogy classes), including the correctly specified theory (i.e., the finest analogy class). Finally, we derive predictions about how the prevalence of analogy-based reasoning may vary with the objective environment: that is, the payoff structure of the stage game.

Consider a centipede game, shown in Figure 2. P1 and P2 take turns choosing Across (A) or Drop (D). The non-terminal nodes are labeled n^k , $1 \leq k \leq K$ where K is an even number. P1 acts at nodes n^1, n^3, \dots, n^{K-1} , P2 acts at nodes n^2, n^4, \dots, n^K , and choosing Drop at n^k leads to the terminal node z^k . If Across is always chosen, then the terminal node z^{end} is reached. If P1 chooses Drop at the first node, the game ends with the payoff profile $(0, 0)$. Every time a player i chooses Across, the sum of payoffs grows by $g > 0$, but if the next player chooses Drop then i 's payoff is $\ell > 0$ smaller than what i would have gotten by choosing Drop. If z^{end} is reached, both get $Kg/2$. That is, if u_j^k is the utility of j at the terminal node z^k , and i moves at n^k , then $u_{-i}^k = u_{-i}^{k-1} - \ell$ while $u_i^k = (u_i^{k-1} + u_{-i}^{k-1} + g) - u_{-i}^k$. This works out to $u_j^k = \frac{g(k-1)}{2}$ for both players when k is odd, and $u_1^k = \frac{k-2}{2}g - \ell$, $u_2^k = \frac{k}{2}g + \ell$ when k is even.

⁷Section 6.2 of Jehiel (2005) mentions that if players could choose their own analogy classes, then the finest analogy classes need not arise, but also says "it is beyond the scope of this paper to analyze the implications of this approach." In a different class of games, Jehiel (1995) similarly observes that another form of bounded rationality (having a limited forecast horizon about opponent's play) can improve welfare.




Figure 2: The centipede game. There are $2K$ non-terminal nodes and players 1 (blue) and 2 (red) alternate in choosing Across (A) or Drop (D). Payoff profiles are shown at the terminal nodes.

While this is an asymmetric stage game, we study the symmetrized version mentioned in Section 2.1, where two matched agents are randomly assigned into the roles of P1 and P2. Let $\mathbb{A} = \{(d^k)_{k=1}^K \in [0, 1]^K\}$, so each strategy is characterized by the probabilities of playing Drop at various nodes in the game tree. When assigned into the role of P1, the strategy (d^k) plays Drop with probabilities d^1, d^3, \dots, d^{K-1} at nodes n^1, n^3, \dots, n^{K-1} . When assigned into the role of P2, it plays Drop with probabilities d^2, d^4, \dots, d^K at nodes n^2, n^4, \dots, n^K . The set of consequences is $\mathbb{Y} = \{1, 2\} \times (\{z_k : 1 \leq k \leq K\} \cup \{z_{end}\})$, where the first dimension of the consequence returns the player role that the agent was assigned into, and the second dimension returns the terminal node reached. The objective distribution over consequences as a function of play is $F^\bullet : \mathbb{A}^2 \rightarrow \Delta(\mathbb{Y})$.

Consider a learning environment where agents know the game tree (i.e., they know F^\bullet), but some agents mistakenly think that when their opponents are assigned into a role, these opponents play Drop with the same probabilities at all of their nodes. Formally, define the restricted space of strategies $\mathbb{A}^{An} := \{(d^k) \in [0, 1]^K : d^k = d^{k'} \text{ if } k \equiv k' \pmod{2}\} \subseteq \mathbb{A}$. The correctly specified extended theory is $\bar{\Theta}^\bullet := \mathbb{A} \times \mathbb{A} \times \{F^\bullet\}$. The misspecified theory with a restriction on beliefs about opponents' play is $\bar{\Theta}^{An} := \mathbb{A}^{An} \times \mathbb{A}^{An} \times \{F^\bullet\}$, reflecting a dogmatic belief that opponents play the same mixed action at all nodes in the analogy class. It is important to remember that these restriction on strategies only exists in the subjective beliefs of the theory $\bar{\Theta}^{An}$ adherents. All agents, regardless of their theory, actually have the strategy space \mathbb{A} . Also, consistent with our framework of theory evolution, agents are born with a theory and cannot choose their own analogy classes.

The next proposition provides a justification for why we might expect agents with coarse analogy classes given by \mathbb{A}^{An} to persist in the society.

Proposition 9. *Suppose $K \geq 4$ and $g > \frac{2}{K-2}\ell$. For any matching assortativity $\lambda \in [0, 1]$, the correctly specified extended theory $\bar{\Theta}^\bullet$ is evolutionarily stable with strategic uncertainty against itself, but it is not evolutionarily stable with strategic uncertainty against the misspecified extended theory $\bar{\Theta}^{An}$. Also, $\bar{\Theta}^{An}$ is not evolutionarily stable against $\bar{\Theta}^\bullet$, unless $\lambda = 1$.*

In contrast to the results from the previous section, which predict different biases may

arise under different matching assortativities, we find in this environment that the correctly specified extended theory is not evolutionarily stable against the extended theory $\bar{\Theta}^{An}$ with coarse analogy classes under any level of assortativity. In the previous application to LQN games, agents with projection bias commit to acting more aggressively, which increases their equilibrium welfare in matches against rational agents but decreases their equilibrium welfare in matches against other agents with the same bias — and vice versa for agents with correlation neglect. But in the current application, the conditional fitness of $\bar{\Theta}^{An}$ against both $\bar{\Theta}^\bullet$ and $\bar{\Theta}^{An}$ can strictly improve on the correctly specified residents' equilibrium fitness. This is because the matches between two adherents of $\bar{\Theta}^\bullet$ must result in Dropping at the first move in equilibrium, while matches where at least one player is an adherent of $\bar{\Theta}^{An}$ either lead to the same outcome or lead to a Pareto dominating payoff profile as the misspecified agent misperceives the opponent's continuation probability and thus chooses Across at almost all of the decision nodes.

But at the same time, $\bar{\Theta}^{An}$ is not evolutionarily stable against $\bar{\Theta}^\bullet$ either. The correctly specified agents can exploit the analogy reasoners' mistake and receive higher payoffs in matches against them, compared with the payoffs that two misspecified agents receive when matched against each other.

This result suggests that neither a homogeneous population of rationals nor a homogeneous population of analogy reasoners can be stable in the long-run, as the resident theory can have lower fitness than the mutants in equilibrium. Our framework makes predictions about the long-run prevalence of analogy-based reasoning in the society, and how such prevalence changes with the stage game. One issue is the multiplicity of EZ-SUs, as there is always the trivial EZ-SU where all agents in both groups always play Drop, and thus learn nothing about others' play. We focus on the EZ-SU where Across is played as often as possible.

Let there be uniform matching in a society with $\Theta_A = \bar{\Theta}^\bullet$, $\Theta_B = \bar{\Theta}^{An}$, population sizes (p_A, p_B) , and suppose $K \geq 4$ and $g > \frac{2}{K-2}\ell$. Consider the *maximal continuation EZ-SU* in the centipede game: this is the strategy profile where $d_{AA}^k = 1$, $d_{BB}^k = 0$ for every $k < K$, $d_{BB}^K = 1$, $d_{AB}^k = 0$ for $k \leq K - 2$, $d_{AB}^{K-1} = d_{AB}^K = 1$, $d_{BA}^k = 0$ for $k < K$, and $d_{BA}^K = 1$. That is, agents always play Across, except the misspecified agents always Drop at node K , the correctly specified agent Drop at nodes $K - 1$ and K when matched with a misspecified agent, and Drop at every node when matched with another correctly specified agent. (The next result verifies that the maximal continuation EZ-SU is indeed an EZ-SU.)

We find the *stable population share* under maximal continuation. This is the unique population size (p_A^*, p_B^*) such that the two theories have the same fitness under the maximal continuation EZ-SU. Intuitively, p_B^* shows what fraction of the society is expected to be analogy reasoners in the long run as the adherents of Θ_A and Θ_B rise and fall according to their fitness, assuming the focal equilibrium is the one with the highest amount of continuation.

Proposition 10. Suppose $K \geq 4$ and $g > \frac{2}{K-2}\ell$. The two theories have the same fitness in the maximal continuation EZ-SU of the centipede game if and only if the population size of the misspecified theory Θ_B is $p_B^* = 1 - \frac{\ell}{g(K-2)}$. This p_B^* is strictly increasing in g and K , strictly decreasing in ℓ .

Under the maintained assumption $g > \frac{2}{K-2}\ell$, the stable population share of misspecified agents is strictly more than 50%, and the share grows when there are more periods and when the increase in payoffs is larger across periods. The main intuition is that the misspecified theory has a higher conditional fitness than the rational theory against rational opponents. The former leads to many periods of continuation and a high payoff for the biased agent when the rational agent eventually drops, but the latter leads to 0 payoff from immediate dropping. The size of the fitness advantage grows with g and K . On the other hand, the misspecified theory has a lower conditional fitness than the rational theory against misspecified opponents. For the two groups to have the same expected fitness, there must be fewer rational opponents (i.e., a smaller stable population share p_A^*) when g and K are higher.

In the centipede game, two successive periods of continuation lead to a strict Pareto improvement in payoffs. Consider instead the dollar game (Reny, 1993) in Figure 3, a variant of the centipede game with a more “competitive” payoff structure, where an agent always gets zero payoff when the opponent plays Drop, no matter how far along in the game tree. Total payoff increases by 1 in each round. If the first player stops immediately, the payoffs are $(1, 0)$, and if the second player continues at the final node n^K , payoffs are $(K + 2, 0)$.




Figure 3: The dollar game. There are K non-terminal nodes and players 1 (blue) and 2 (red) alternate in choosing Across (A) or Drop (D). Payoff profiles are shown at the terminal nodes.

Proposition 11. For every population size $(p, 1 - p)$ with $p \in [0, 1]$, the maximal continuation EZ-SU is an EZ-SU where the fitness of $\bar{\Theta}^\bullet$ is strictly higher than that of $\bar{\Theta}^{An}$.

While the maximal continuation EZ-SU remains an EZ-SU in the dollar game, the rational theory always has strictly higher fitness than the misspecified theory with analogy classes with any population sizes. So, provided the maximal continuation EZ-SU remains focal, then we should expect no analogy reasoners in the long run with this different stage game. This is because the misspecified theory now has a lower conditional fitness than the rational theory

against rational opponents. A misspecified P1 gets 0 in the maximal continuation EZ-SU against a rational P2 who drops at their penultimate decision node, whereas a rational P1 gets 1 from immediate dropping.

In a recent survey article on ABEEs, Jehiel (2020) points out that the misspecified Bayesian learning approach to analogy classes should aim for “a better understanding of how the subjective theories considered by the players may be shaped by the objective characteristics of the environment.”⁸ We find that a framework of theory evolution can deliver such results. Taken together, the analyses of maximal continuation EZ-SUs in the centipede game and the dollar game suggest that analogy-based reasoning should be more prevalent in the centipede game than the dollar game because the payoff structure is “less competitive” in the former. Also, in the centipede game, the prevalence of this bias should increase with a longer horizon and with faster payoff growth.

6 Concluding Discussion

This paper presents an evolutionary selection criterion to endogenize (mis)specifications when agents learn about a strategic environment. We introduce the concept of a zeitgeist to capture the ambient social structure where learning takes place: the prominence of different theories in the society and the interaction patterns among their adherents. These details matter because different types of opponents behave differently, inducing different beliefs about the economic fundamentals for a misspecified agent. Evolutionary stability of a theory is defined based on the expected objective payoffs (fitness) of its adherents in equilibrium.

We have highlighted settings where the correct specification is not evolutionarily stable against some misspecifications. We view our main contributions as two fold. First, we point out how details of the zeitgeist (e.g., the matching assortativity) change which learning biases may persist in an otherwise rational society. Second, we emphasize that the learning channel, unique to a world where evolutionary forces act on specifications (sets of feasible beliefs) instead of single beliefs, generates novel stability phenomena.

Our framework evaluates whether a misspecification is likely to persist once it emerges in a society, but does not account for which errors appear in the first place. It is plausible that some first-stage filter prevents certain obvious misspecifications from ever reaching the stage that we study in the evolutionary framework. In the applications, we have focused on

⁸Jehiel (2020) interprets ABEEs as players adopting the “simplest” explanations of observed aggregate statistics of play in environments with physically or psychologically coarse feedback. An objectively coarse feedback structure can lead agents to adopt the subjective theory that others behave in the same way in all contingencies in the same coarse analogy class. Our approach suggests that the very prevalence of analogy classes may also be shaped by other objective environmental primitives, including the payoff structure of the stage game.

misspecifications that seem psychologically plausible or harder to detect, such as misspecified higher-order beliefs.

We have used the simplest evolutionary framework where fitness is identified with the expectation of objective payoffs, as opposed to some more exotic function of the payoffs. This paper not meant to be a just-so congruence exercise of identifying the suitable definition of fitness to justify a particular error (which is the focus for many of the papers that [Robson and Samuelson \(2011\)](#) survey). Rather, we hope that our stability notions are reasonably simple and universal that they may become a part of the applied theory toolkit in the future. Studies on the implications of misspecifications in various strategic environments may further enrich our understanding of these errors by paying more attention to their evolutionary stability.

References

- ALGER, I. AND J. WEIBULL (2019): “Evolutionary models of preference formation,” *Annual Review of Economics*, 11, 329–354.
- ALIPRANTIS, C. AND K. BORDER (2006): *Infinite Dimensional Analysis: A Hitchhiker’s Guide*, Springer Science & Business Media.
- ANGELETOS, G.-M. AND A. PAVAN (2007): “Efficient use of information and social value of information,” *Econometrica*, 75, 1103–1142.
- BA, C. (2020): “Model misspecification and paradigm shift,” *Working Paper*.
- BA, C. AND A. GINDIN (2020): “A multi-agent model of misspecified learning with overconfidence,” *Working Paper*.
- BERGEMANN, D. AND S. MORRIS (2013): “Robust predictions in games with incomplete information,” *Econometrica*, 81, 1251–1308.
- BERMAN, R. AND Y. HELLER (2020): “Naive analytics equilibrium,” *Working Paper*.
- BLUME, L. AND D. EASLEY (2002): “Optimality and natural selection in markets,” *Journal of Economic Theory*, 107, 95–135.
- (2006): “If you’re so smart, why aren’t you rich? Belief selection in complete and incomplete markets,” *Econometrica*, 74, 929–966.
- BOHREN, J. A. (2016): “Informational herding with model misspecification,” *Journal of Economic Theory*, 163, 222–247.
- BOHREN, J. A. AND D. HAUSER (2018): “Learning with model misspecification: Characterization and robustness,” *Working Paper*.

- CHO, I.-K. AND K. KASA (2015): “Learning and model validation,” *Review of Economic Studies*, 82, 45–82.
- (2017): “Gresham’s law of model averaging,” *American Economic Review*, 107, 3589–3616.
- DASARATHA, K. AND K. HE (2020): “Network structure and naive sequential learning,” *Theoretical Economics*, 15, 415–444.
- DEKEL, E., J. ELY, AND O. YILANKAYA (2007): “Evolution of preferences,” *Review of Economic Studies*, 74, 685–704.
- ELIAZ, K. AND R. SPIEGLER (2020): “A model of competing narratives,” *American Economic Review*, 110, 3786–3816.
- 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,” *Working Paper*.
- FRICK, M., R. IJIMA, AND Y. ISHII (2019): “Stability and robustness in misspecified learning models,” *Working Paper*.
- (2020): “Misinterpreting others and the fragility of social learning,” *Econometrica*, 88, 2281–2328.
- (2021): “Welfare comparisons for biased learning,” *Working Paper*.
- FRIEDMAN, M. (1953): *Essays in Positive Economics*, University of Chicago Press.
- FUDENBERG, D. AND G. LANZANI (2020): “Which misperceptions persist?” *Working Paper*.
- FUDENBERG, D., G. LANZANI, AND P. STRACK (2020): “Limits points of endogenous misspecified learning,” *Working Paper*.
- FUDENBERG, D., G. ROMANYUK, AND P. STRACK (2017): “Active learning with a misspecified prior,” *Theoretical Economics*, 12, 1155–1189.
- GAGNON-BARTSCH, T., M. RABIN, AND J. SCHWARTZSTEIN (2020): “Channeled attention and stable errors,” *Working Paper*.
- GOLDFARB, A. AND M. XIAO (2019): “Transitory shocks, limited attention, and a firm’s decision to exit,” *Working Paper*.
- HANSEN, K., K. MISRA, AND M. PAI (2021): “Algorithmic collusion: Supra-competitive prices via independent algorithms,” *Marketing Science*, forthcoming.
- HE, K. (2020): “Mislearning from censored data: The gambler’s fallacy in optimal-stopping problems,” *Working Paper*.

- HEIDHUES, P., B. KOSZEGI, AND P. STRACK (2018): “Unrealistic expectations and misguided learning,” *Econometrica*, 86, 1159–1214.
- HELLER, Y. (2015): “Three steps ahead,” *Theoretical Economics*, 10, 203–241.
- HELLER, Y. AND E. WINTER (2016): “Rule rationality,” *International Economic Review*, 57, 997–1026.
- (2020): “Biased-belief equilibrium,” *American Economic Journal: Microeconomics*, 12, 1–40.
- JEHIEL, P. (1995): “Limited horizon forecast in repeated alternate games,” *Journal of Economic Theory*, 67, 497–519.
- (2005): “Analogy-based expectation equilibrium,” *Journal of Economic theory*, 123, 81–104.
- (2018): “Investment strategy and selection bias: An equilibrium perspective on overoptimism,” *American Economic Review*, 108, 1582–97.
- (2020): “Analogy-based expectation equilibrium and related concepts: Theory, applications, and beyond,” *Working Paper*.
- LEVY, G., R. RAZIN, AND A. YOUNG (2020): “Misspecified politics and the recurrence of populism,” *Working Paper*.
- MASSARI, F. (2020): “Under-reaction: Irrational behavior or robust response to model misspecification?” *Working Paper*.
- MOLAVI, P. (2019): “Macroeconomics with learning and misspecification: A general theory and applications,” *Working Paper*.
- MUROOKA, T. AND Y. YAMAMOTO (2021): “Multi-Player Bayesian Learning with Misspecified Models,” *Working Paper*.
- 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 (2020): “Competing models,” *Working Paper*.
- RENY, P. J. (1993): “Common belief and the theory of games with perfect information,” *Journal of Economic Theory*, 59, 257–274.
- ROBSON, A. J. AND L. SAMUELSON (2011): “The evolutionary foundations of preferences,” in *Handbook of Social Economics*, Elsevier, vol. 1, 221–310.
- SANDRONI, A. (2000): “Do markets favor agents able to make accurate predictions?” *Econometrica*, 68, 1303–1341.

SCHWARTZSTEIN, J. AND A. SUNDERAM (2021): “Using models to persuade,” *American Economic Review*, *forthcoming*.

VIVES, X. (1988): “Aggregation of information in large Cournot markets,” *Econometrica*, 851–876.

Appendix

A Proofs of Results from the Main Text

A.1 Proof of Proposition 1

Proof. In any approachable EZ, let $F \in \text{supp}(\mu_A)$ and note that $F^\bullet \in \Theta_A$ since Θ_A is correctly specified. Both F and F^\bullet solve the weighted minimization problem, the former because it is in the support of μ_A , the latter because it attains the lowest minimization objective of 0. By strong identification, the set of best responses to a_{AA} and a_{BA} under the belief μ_A is the same as set of actions that maximize payoffs in the decision problem given by F^\bullet . Therefore, adherents of Θ_A obtain the highest possible objective payoffs in the stage game, so Θ_A has weakly higher fitness than Θ_B in the approachable EZ. \square

A.2 Proof of Proposition 2

Proof. Let two singleton theories Θ_A, Θ_B be given. By way of contradiction, suppose they exhibit stability reversal. Let $\mathfrak{Z} = (\Theta_A, \Theta_B, \mu_A, \mu_B, p = (0, 1), \lambda = 0, (a))$ be any EZ where Θ_B is resident. By the definition of EZ, $\mathfrak{Z}' = (\Theta_A, \Theta_B, \mu_A, \mu_B, p = (1, 0), \lambda = 0, (a))$ is also an EZ where Θ_A is resident. Let $u_{g,g'}$ be theory Θ_g 's conditional fitness against group g' in the EZ \mathfrak{Z}' . Part (i) of the definition of stability reversal requires that $u_{AA} > u_{BA}$ and $u_{AB} > u_{BB}$. These conditional fitness levels remain the same in \mathfrak{Z} . This means the fitness of Θ_A is strictly higher than that of Θ_B in \mathfrak{Z} , a contradiction. \square

A.3 Proof of Example 1

Proof. Define $b^*(a_i, a_{-i}) := b^\bullet + \frac{m}{a_i + a_{-i}}$. It is clear that $D_{KL}(F^\bullet(a_i, a_{-i}) \parallel \hat{F}(a_i, a_{-i}; b^*(a_i, a_{-i}), m)) = 0$, while this KL divergence is strictly positive for any other choice of b .

In every EZ with $\lambda = 0$ and $p = (1, 0)$, we must have $a_{AA} = a_{AB} = 1$. If $a_{BA} = 2$, then the adherents of Θ_B infer $b^*(1, 2) = b^\bullet + \frac{m}{3}$. With this inference, the biased agents expect $1 \cdot (2(b^\bullet + \frac{m}{3}) - m) = 2b^\bullet - \frac{m}{3}$ from playing 1 against rival investment 1, and expect $2 \cdot (3(b^\bullet + \frac{m}{3}) - m) - c = 6b^\bullet - c$ from playing 2 against rival investment 1. Since $4b^\bullet + \frac{m}{3} - c > 0$ from Condition 2, there is an EZ with $a_{BA} = 2$ and μ_B puts probability 1 on $b^\bullet + \frac{m}{3}$. It is impossible to have $a_{BA} = 1$ in EZ. This is because $b^*(1, 1) > b^*(1, 2)$, and under the inference $b^*(1, 2)$ we already have that the best response to 1 is 2, so the same also holds under any higher belief about complementarity. Also, we have $a_{BB} = 2$, since 2 must best respond to both 1 and 2. So in every such EZ, Θ_A 's conditional fitness against group A is $2b^\bullet$ and Θ_B 's conditional fitness against group A is $6b^\bullet - c$, with $2b^\bullet > 6b^\bullet - c$ by Condition 1. Also, Θ_A 's

conditional fitness against group B is $3b^\bullet$, while Θ_B 's conditional fitness against group B is $8b^\bullet - c$. Again, $3b^\bullet > 8b^\bullet - c$ by Condition 1.

Next, we show Θ_B has strictly higher fitness than Θ_A in every EZ with $\lambda = 0, p_B = 1$. There is no EZ with $a_{BB} = 1$. This is because $b^*(1, 1) = b^\bullet + \frac{m}{2}$. As discussed before, under this inference the best response to 1 is 2, not 1. Now suppose $a_{BB} = 2$. Then μ_B puts probability 1 on $b^*(2, 2) = b^\bullet + \frac{m}{4}$. With this inference, the biased agents expect $1 \cdot (3(b^\bullet + \frac{m}{4}) - m) = 3b^\bullet - \frac{m}{4}$ from playing 1 against rival investment 2, and expect $2 \cdot (4(b^\bullet + \frac{m}{4}) - m) - c = 8b^\bullet - c$ from playing 2 against rival investment 2. We have $5b^\bullet + \frac{m}{4} - c > 0$ from Condition 2, so 2 best responds to 2. We must have $a_{AA} = a_{AB} = 1$. We conclude the unique EZ behavior is $(a_{AA}, a_{AB}, a_{BA}, a_{BB}) = (1, 1, 1, 2)$, since the biased agents expect $1 \cdot (2(b^\bullet + \frac{m}{4}) - m) = 2b^\bullet - \frac{m}{2}$ from playing 1 against rival investment 1, and expect $2 \cdot (3(b^\bullet + \frac{m}{4}) - m) - c = 6b^\bullet - \frac{m}{2} - c$ from playing 2 against rival investment 1. We have $4b^\bullet - c < 0$ from Condition 1, so 1 best responds to 1. In the unique EZ with $\lambda = 0$ and $p = (0, 1)$, the fitness of Θ_A is $2b^\bullet$ and the fitness of Θ_B is $8b^\bullet - c$, where $8b^\bullet - c > 2b^\bullet$ by Condition 1. \square

A.4 Proof of Proposition 3

Proof. To show the first claim, by way of contradiction, suppose $\mathfrak{Z} = (\Theta_A, \Theta_B, \mu_A, \mu_B, p = (1, 0), \lambda = 0, (a_{AA}, a_{AB}, a_{BA}, a_{BB}))$ is an EZ, and $\tilde{\mathfrak{Z}} = (\Theta_A, \Theta_B, \mu_A, \mu_B, p = (0, 1), \lambda = 0, (\tilde{a}_{AA}, \tilde{a}_{AB}, \tilde{a}_{BA}, \tilde{a}_{BB}))$ is another EZ where the adherents of Θ_B hold the same belief μ_B (group A's belief cannot change as Θ_A is the correctly specified singleton theory). By the optimality of behavior in \mathfrak{Z} , a_{BA} best responds to a_{AB} under the belief μ_B , and a_{AB} best responds to a_{BA} under the belief μ_A , therefore $\tilde{\mathfrak{Z}}' = (\Theta_A, \Theta_B, \mu_A, \mu_B, p = (0, 1), \lambda = 0, (\tilde{a}_{AA}, a_{AB}, a_{BA}, \tilde{a}_{BB}))$ is another EZ. This holds because the distributions of observations for the adherents of Θ_B are identical in $\tilde{\mathfrak{Z}}$ and $\tilde{\mathfrak{Z}}'$, since they only face data generated from the profile $(\tilde{a}_{BB}, \tilde{a}_{BB})$. At the same time, since \tilde{a}_{BB} best responds to itself under the belief μ_B , we have that $\mathfrak{Z}' = (\Theta_A, \Theta_B, \mu_A, \mu_B, p = (1, 0), \lambda = 0, (a_{AA}, a_{AB}, a_{BA}, \tilde{a}_{BB}))$ is an EZ. Part (i) of the definition of stability reversal applied to \mathfrak{Z}' requires that $U^\bullet(a_{AB}, a_{BA}) > U^\bullet(\tilde{a}_{BB}, \tilde{a}_{BB})$ (where U^\bullet is the objective expected payoffs), but part (ii) of the same definition applied to $\tilde{\mathfrak{Z}}'$ requires $U^\bullet(\tilde{a}_{BB}, \tilde{a}_{BB}) \geq U^\bullet(a_{AB}, a_{BA})$, a contradiction.

To show the second claim, by way of contradiction suppose Θ_B is strategically independent and $\mathfrak{Z} = (\Theta_A, \Theta_B, \mu_A, \mu_B, p = (0, 1), \lambda = 0, (a_{AA}, a_{AB}, a_{BA}, a_{BB}))$ is an EZ. By strategic independence, the adherents of Θ_B find it optimal to play a_{BB} against any opponent strategy under the belief μ_B . So, there exists another EZ of the form $\mathfrak{Z}' = (\Theta_A, \Theta_B, \mu'_A, \mu_B, p = (0, 1), \lambda = 0, (a_{AA}, a'_{AB}, a_{BB}, a_{BB}))$, where a'_{AB} is an objective best response to a_{BB} . The belief μ_B is sustained because in both \mathfrak{Z} and \mathfrak{Z}' , the adherents of Θ_B have the same data: from the strategy profile (a_{BB}, a_{BB}) . In \mathfrak{Z}' , Θ_A 's fitness is $U^\bullet(a'_{AB}, a_{BB})$ and Θ_B 's fitness is

$U^\bullet(a_{BB}, a_{BB})$. We have $U^\bullet(a'_{AB}, a_{BB}) \geq U^\bullet(a_{BB}, a_{BB})$ since a'_{AB} is an objective best response to a_{BB} , contradicting the definition of stability reversal. \square

A.5 Proof of Proposition 4

Proof. Let $\lambda \in [0, 1]$ be given and let $\mathfrak{Z} = (\Theta_A, \Theta_B, \mu_A, \mu_B, p = (1, 0), \lambda, (a))$ be an EZ. Since Θ_A, Θ_B are singleton theories, $\mathfrak{Z}_0 = (\Theta_A, \Theta_B, \mu_A, \mu_B, p = (1, 0), \lambda = 0, (a))$ and $\mathfrak{Z}_1 = (\Theta_A, \Theta_B, \mu_A, \mu_B, p = (1, 0), \lambda = 1, (a))$ are also EZs. Furthermore, they are all approachable since the same beliefs and behavior are sustained as EZs with any population proportions. Let $u_{g,g'}$ represent theory Θ_g 's conditional fitness against group g' in each of these three EZs. From the hypothesis of the proposition, $u_{A,A} \geq u_{B,A}$ and $u_{A,A} \geq u_{B,B}$. This means the fitness of Θ_A in \mathfrak{Z} , which is $u_{A,A}$, is weakly larger than the fitness of Θ_B in \mathfrak{Z} , which is $\lambda u_{B,B} + (1 - \lambda)u_{B,A}$. This shows Θ_A has weakly higher fitness than Θ_B in every approachable EZ with λ and $p = (1, 0)$. Also, at least one such approachable EZ exists with assortativity λ , for at least one approachable EZ exists when $\lambda = 0$, and the same equilibrium belief and behavior also constitutes an EZ for any other assortativity. \square

A.6 Proof of Example 2

Proof. Let $KL_{4,1} := 0.4 \cdot \ln \frac{0.4}{0.1} + 0.6 \cdot \ln \frac{0.6}{0.9} \approx 0.3112$, $KL_{4,8} := 0.4 \cdot \ln \frac{0.4}{0.8} + 0.6 \cdot \ln \frac{0.6}{0.2} \approx 0.3819$, and $KL_{2,4} := 0.2 \cdot \ln \frac{0.2}{0.4} + 0.8 \cdot \ln \frac{0.8}{0.6} \approx 0.0915$. Let λ_h be the unique solution to $(1 - \lambda)KL_{2,4} - \lambda(KL_{4,8} - KL_{4,1}) = 0$, so $\lambda_h \approx 0.564$.

We show for any $\lambda \in [0, \lambda_h)$, there exists a unique EZ $\mathfrak{Z} = (\Theta_A, \Theta_B, \mu_A, \mu_B, p = (1, 0), \lambda, (a))$, and that this EZ has μ_B putting probability 1 on F_H , $a_{AA} = a_1$, $a_{AB} = a_1$, $a_{BA} = a_2$, $a_{BB} = a_2$. First, we may verify that under F_H , a_2 best responds to both a_1 and a_2 . Also, the KL divergence of F_H is $\lambda \cdot KL_{4,8}$ while that of F_L is $\lambda \cdot KL_{4,1} + (1 - \lambda) \cdot KL_{2,4}$. Since $\lambda < \lambda_h$, we see that F_H has strictly lower KL divergence. Finally, to check that there are no other EZs, note we must have $a_{AA} = a_1$, $a_{AB} = a_1$, $a_{BA} = a_2$ in every EZ. In an EZ where a_{BB} puts probability $q \in [0, 1]$ on a_2 , the KL divergence of F_H is $\lambda p \cdot KL_{4,8}$ and the KL divergence of F_L is $\lambda p \cdot KL_{4,1} + (1 - \lambda) \cdot KL_{2,4}$. We have

$$\lambda q \cdot KL_{4,1} + (1 - \lambda) \cdot KL_{2,4} - \lambda q \cdot KL_{4,8} = \lambda q \cdot (KL_{4,1} - KL_{4,8}) + (1 - \lambda) KL_{2,4} \geq (1 - \lambda) KL_{2,4} - \lambda (KL_{4,8} - KL_{4,1}).$$

Since $\lambda < \lambda_h$, this is strictly positive. Therefore we must have μ_B put probability 1 on F_H , which in turn implies $q = 1$.

For each $\lambda \in [0, \lambda_h)$, the beliefs and behavior in the unique EZ discussed above also constitute an EZ for a small enough $p_B > 0$. So, the unique EZ with $p_B = 0$ is approachable.

When Θ_A is dominant, the equilibrium fitness of Θ_A is always 0.25 for every λ . The

equilibrium fitness of Θ_B , as a function of λ , is $0.4\lambda + 0.2(1 - \lambda)$. Let λ_l solve $0.25 = 0.4\lambda + 0.2(1 - \lambda)$, that is $\lambda_l = 0.25$. This shows Θ_A is evolutionarily fragile against Θ_B for $\lambda \in (\lambda_l, \lambda_h)$, and it is evolutionarily stable against Θ_B for $\lambda = 0$.

Now suppose $\lambda = 1$. If there is an EZ with $p_A = 1$ where a_{BB} plays a_2 with positive probability, then μ_B must put probability 1 on F_L , since $KL_{4,1} < KL_{4,8}$. This is a contradiction, since a_2 does not best respond to itself under F_L . So the unique EZ involves $a_{AA} = a_1$, $a_{AB} = a_1$, $a_{BA} = a_2$, $a_{BB} = a_3$. It is easy to check this EZ is approachable. In the EZ, the fitness of Θ_A is 0.25, and the fitness of Θ_B is 0.2. This shows Θ_A is evolutionarily stable against Θ_B for $\lambda = 1$. \square

A.7 Proof of Lemma 1

Proof. For $i \neq j$, rewrite $s_i = \left(\omega + \frac{\kappa}{\sqrt{\kappa^2 + (1-\kappa)^2}}z\right) + \frac{1-\kappa}{\sqrt{\kappa^2 + (1-\kappa)^2}}\eta_i$ and $s_j = \left(\omega + \frac{\kappa}{\sqrt{\kappa^2 + (1-\kappa)^2}}z\right) + \frac{1-\kappa}{\sqrt{\kappa^2 + (1-\kappa)^2}}\eta_j$. Note that $\omega + \frac{\kappa}{\sqrt{\kappa^2 + (1-\kappa)^2}}z$ has a normal distribution with mean 0 and variance $\sigma_\omega^2 + \frac{\kappa^2}{\kappa^2 + (1-\kappa)^2}\sigma_\epsilon^2$. The posterior distribution of $\left(\omega + \frac{\kappa}{\sqrt{\kappa^2 + (1-\kappa)^2}}z\right)$ given s_i is therefore normal

with a mean of $\frac{1/(\frac{(1-\kappa)^2}{\kappa^2 + (1-\kappa)^2}\sigma_\epsilon^2)}{1/(\sigma_\omega^2 + \frac{\kappa^2}{\kappa^2 + (1-\kappa)^2}\sigma_\epsilon^2) + 1/(\frac{(1-\kappa)^2}{\kappa^2 + (1-\kappa)^2}\sigma_\epsilon^2)}s_i$ and a variance of $\frac{1}{1/(\sigma_\omega^2 + \frac{\kappa^2}{\kappa^2 + (1-\kappa)^2}\sigma_\epsilon^2) + 1/(\frac{(1-\kappa)^2}{\kappa^2 + (1-\kappa)^2}\sigma_\epsilon^2)}$.

Since η_j is mean-zero and independent of i 's signal, the posterior distribution of $s_j \mid s_i$ under the correlation parameter κ is normal with a mean of

$$\frac{1/(\frac{(1-\kappa)^2}{\kappa^2 + (1-\kappa)^2}\sigma_\epsilon^2)}{1/(\sigma_\omega^2 + \frac{\kappa^2}{\kappa^2 + (1-\kappa)^2}\sigma_\epsilon^2) + 1/(\frac{(1-\kappa)^2}{\kappa^2 + (1-\kappa)^2}\sigma_\epsilon^2)}s_i$$

and a variance of $\frac{1}{1/(\sigma_\omega^2 + \frac{\kappa^2}{\kappa^2 + (1-\kappa)^2}\sigma_\epsilon^2) + 1/(\frac{(1-\kappa)^2}{\kappa^2 + (1-\kappa)^2}\sigma_\epsilon^2)} + \frac{(1-\kappa)^2}{\kappa^2 + (1-\kappa)^2}\sigma_\epsilon^2$. We thus define $\psi(\kappa) :=$

$\frac{1/(\frac{(1-\kappa)^2}{\kappa^2 + (1-\kappa)^2}\sigma_\epsilon^2)}{1/(\sigma_\omega^2 + \frac{\kappa^2}{\kappa^2 + (1-\kappa)^2}\sigma_\epsilon^2) + 1/(\frac{(1-\kappa)^2}{\kappa^2 + (1-\kappa)^2}\sigma_\epsilon^2)}$ for $\kappa \in [0, 1)$, and $\psi(1) := 1$. To see that $\psi(\kappa)$ is strictly increasing in k , we have

$$\begin{aligned} 1/\psi(\kappa) &= 1 + \frac{\frac{(1-\kappa)^2}{\kappa^2 + (1-\kappa)^2}\sigma_\epsilon^2}{\sigma_\omega^2 + \frac{\kappa^2}{\kappa^2 + (1-\kappa)^2}\sigma_\epsilon^2} \\ &= 1 + \frac{(1-\kappa)^2\sigma_\epsilon^2}{(\kappa^2 + (1-\kappa)^2)\sigma_\omega^2 + \kappa^2\sigma_\epsilon^2} \end{aligned}$$

and then we can verify that the second term is decreasing in κ .

As $\kappa \rightarrow 1$, the term $1/(\frac{(1-\kappa)^2}{\kappa^2 + (1-\kappa)^2}\sigma_\epsilon^2)$ tends to ∞ , so $\frac{1/(\frac{(1-\kappa)^2}{\kappa^2 + (1-\kappa)^2}\sigma_\epsilon^2)}{1/(\sigma_\omega^2 + \frac{\kappa^2}{\kappa^2 + (1-\kappa)^2}\sigma_\epsilon^2) + 1/(\frac{(1-\kappa)^2}{\kappa^2 + (1-\kappa)^2}\sigma_\epsilon^2)}$ approaches $\frac{1/(\frac{(1-\kappa)^2}{\kappa^2 + (1-\kappa)^2}\sigma_\epsilon^2)}{1/(\frac{(1-\kappa)^2}{\kappa^2 + (1-\kappa)^2}\sigma_\epsilon^2)} = 1$. We also verify that $\psi(0) = \frac{1/\sigma_\epsilon^2}{(1/\sigma_\omega^2) + (1/\sigma_\epsilon^2)} > 0$.

Finally, for any $\kappa \in [0, 1]$, $\frac{\kappa}{\sqrt{\kappa^2+(1-\kappa)^2}}z + \frac{1-\kappa}{\sqrt{\kappa^2+(1-\kappa)^2}}\eta_i$ has variance σ_ϵ^2 and mean 0, so $\mathbb{E}_\kappa[\omega \mid s_i] = \frac{1/\sigma_\epsilon^2}{1/\sigma_\epsilon^2+1/\sigma_\omega^2}s_i$. We then define γ as the strictly positive constant $\frac{1/\sigma_\epsilon^2}{1/\sigma_\epsilon^2+1/\sigma_\omega^2}$. \square

A.8 Proof of Lemma 2

Proof. Player i 's conditional expected utility given signal s_i is

$$\alpha_i s_i \cdot \mathbb{E}_\kappa[\mathbb{E}_{r \sim \text{marg}_r(\mu)}[\omega - \frac{1}{2}r\alpha_i s_i - \frac{1}{2}r\alpha_{-i}s_{-i} + \zeta \mid s_i] - \frac{1}{2}(\alpha_i s_i)^2]$$

by linearity, expectation over r is equivalent to evaluating the inner expectation with $r = \hat{r}$, which gives

$$\begin{aligned} & \alpha_i s_i \cdot \mathbb{E}_\kappa[\omega - \frac{1}{2}\hat{r}\alpha_i s_i - \frac{1}{2}\hat{r}\alpha_{-i}s_{-i} + \zeta \mid s_i] - \frac{1}{2}(\alpha_i s_i)^2 \\ &= \alpha_i s_i \cdot (\gamma s_i - \frac{1}{2}\hat{r}\alpha_i s_i - \frac{1}{2}\hat{r}\psi(\kappa)s_i\alpha_{-i}) - \frac{1}{2}(\alpha_i s_i)^2 \\ &= s_i^2 \cdot (\alpha_i \gamma - \frac{1}{2}\hat{r}\alpha_i^2 - \frac{1}{2}\hat{r}\psi(\kappa)\alpha_i\alpha_{-i} - \frac{1}{2}\alpha_i^2). \end{aligned}$$

The term in parenthesis does not depend on s_i , and the second moment of s_i is the same for all values of κ . Therefore this expectation is $\mathbb{E}[s_i^2] \cdot (\alpha_i \gamma - \frac{1}{2}\hat{r}\alpha_i^2 - \frac{1}{2}\hat{r}\psi(\kappa)\alpha_i\alpha_{-i} - \frac{1}{2}\alpha_i^2)$. The expression for $\alpha_i^{BR}(\alpha_{-i}; \kappa, r)$ follows from simple algebra, noting that $\mathbb{E}[s_i^2] > 0$ while the second derivative with respect to α_i for the term in the parenthesis is $-\frac{1}{2}\hat{r} - \frac{1}{2} < 0$.

To see that the said linear strategy is optimal among all strategies, suppose i instead chooses any q_i after s_i . By above arguments, the objective to maximize is

$$q_i \cdot (\gamma s_i - \frac{1}{2}\hat{r}q_i - \frac{1}{2}\hat{r}\psi(\kappa)s_i\alpha_{-i}) - \frac{1}{2}q_i^2.$$

This objective is a strictly concave function in q_i , as $-\frac{1}{2}\hat{r} - \frac{1}{2} < 0$. First-order condition finds the maximizer $q_i^* = \alpha_i^{BR}(\alpha_{-i}; \kappa, \hat{r})$. Therefore, the linear strategy also maximizes interim expected utility after every signal s_i , and so it cannot be improved on by any other strategy. \square

A.9 Proof of Lemma 3

Proof. Note that $\frac{\alpha_i + \alpha_{-i}\psi(\kappa^\bullet)}{\alpha_i + \alpha_{-i}\psi(\kappa)} \geq 0$ and $\frac{\alpha_i + \alpha_{-i}\psi(\kappa^\bullet)}{\alpha_i + \alpha_{-i}\psi(\kappa)} = 1 + \frac{\alpha_{-i}(\psi(\kappa^\bullet) - \psi(\kappa))}{\alpha_i + \alpha_{-i}\psi(\kappa)} \leq 1 + \frac{1}{\psi(0)}$ (recalling $\psi(0) > 0$). Hence let $L_3 = r^\bullet \cdot (1 + \frac{1}{\psi(0)})$. When $\bar{M}_r \geq L_3$, we always have $r_i^{INF}(\alpha_i, \alpha_{-i}; \kappa^\bullet, \kappa, r^\bullet) \leq \bar{M}_r$ for all $\alpha_i, \alpha_{-i} \geq 0$ and $\kappa^\bullet, \kappa \in [0, 1]$.

Conditional on the signal s_i , the distribution of market price under the model $F_{\hat{r}, \kappa, \hat{\sigma}_\zeta}$ is

normal with a mean of

$$\mathbb{E}[\omega \mid s_i] - \frac{1}{2}\hat{r}\alpha_i s_i - \frac{1}{2}\hat{r}\alpha_{-i} \cdot \mathbb{E}_{\kappa}[s_{-i} \mid s_i] = \gamma s_i - \frac{1}{2}\hat{r}\alpha_i s_i - \frac{1}{2}\hat{r}\alpha_{-i}\psi(\kappa)s_i,$$

while the distribution of market price under the model $F_{r^\bullet, \kappa^\bullet, \sigma_\zeta^\bullet}$ is normal with a mean of

$$\mathbb{E}[\omega \mid s_i] - \frac{1}{2}r^\bullet\alpha_i s_i - \frac{1}{2}r^\bullet\alpha_{-i} \cdot \mathbb{E}_{\kappa^\bullet}[s_{-i} \mid s_i] = \gamma s_i - \frac{1}{2}r^\bullet\alpha_i s_i - \frac{1}{2}r^\bullet\alpha_{-i}\psi(\kappa^\bullet)s_i.$$

Matching coefficients on s_i , we find that if $\hat{r} = r^\bullet \frac{\alpha_i + \alpha_{-i}\psi(\kappa^\bullet)}{\alpha_i + \alpha_{-i}\psi(\kappa)}$, then these means match after every s_i . On the other hand, for any other value of \hat{r} , these means will not match for any s_i and thus $D_{KL}(F_{r^\bullet, \kappa^\bullet, \sigma_\zeta^\bullet}(\alpha_i, \alpha_{-i}) \parallel F_{\hat{r}, \kappa, \hat{\sigma}_\zeta}(\alpha_i, \alpha_{-i})) > 0$ for any $\hat{r} \neq r^\bullet \frac{\alpha_i + \alpha_{-i}\psi(\kappa^\bullet)}{\alpha_i + \alpha_{-i}\psi(\kappa)}$.

Let $L_1 = \max_{\kappa \in [0,1]} \left\{ \text{Var}_\kappa[\omega \mid s_i] + \text{Var}_\kappa \left[\frac{1}{2}r^\bullet \cdot \left(1 + \frac{1}{\psi(0)}\right) B_\alpha \cdot s_{-i} \mid s_i \right] \right\}$. This maximum exists and is finite, since the expression is a continuous function of κ on the compact domain $[0, 1]$. Also, let $L_2 = \max_{\kappa \in [0,1]} \left\{ \text{Var}_\kappa[\omega \mid s_i] + \text{Var}_\kappa \left[\frac{1}{2}r^\bullet B_\alpha \cdot s_{-i} \mid s_i \right] \right\}$, where the maximum exists for the same reason. Conditional on the signal s_i , the variance of market price under the model $F_{r^\bullet \frac{\alpha_i + \alpha_{-i}\psi(\kappa^\bullet)}{\alpha_i + \alpha_{-i}\psi(\kappa)}, \kappa, \hat{\sigma}_\zeta}$ is

$$\text{Var}_\kappa \left[\omega - \frac{1}{2}r^\bullet \frac{\alpha_i + \alpha_{-i}\psi(\kappa^\bullet)}{\alpha_i + \alpha_{-i}\psi(\kappa)} \alpha_{-i} s_{-i} \mid s_i \right] + \hat{\sigma}_\zeta^2.$$

Since ω and s_{-i} are positively correlated given s_i , and using the fact $r^\bullet \frac{\alpha_i + \alpha_{-i}\psi(\kappa^\bullet)}{\alpha_i + \alpha_{-i}\psi(\kappa)} \leq r^\bullet \cdot \left(1 + \frac{1}{\psi(0)}\right)$ and $\alpha_{-i} \leq B_\alpha$, this variance is no larger than

$$\text{Var}_\kappa [\omega \mid s_i] + \text{Var}_\kappa \left[\frac{1}{2}r^\bullet \cdot \left(1 + \frac{1}{\psi(0)}\right) B_\alpha \cdot s_{-i} \mid s_i \right] + \hat{\sigma}_\zeta^2 = L_1 + \hat{\sigma}_\zeta^2.$$

On the other hand, the variance of market price under the model $F_{r^\bullet, \kappa^\bullet, \sigma_\zeta^\bullet}$ is

$$\text{Var}_{\kappa^\bullet} \left[\omega - \frac{1}{2}r^\bullet \alpha_{-i} s_{-i} \mid s_i \right] + (\sigma_\zeta^\bullet)^2 \leq \text{Var}_{\kappa^\bullet} [\omega \mid s_i] + \text{Var}_{\kappa^\bullet} \left[\frac{1}{2}r^\bullet B_\alpha \cdot s_{-i} \mid s_i \right] + (\sigma_\zeta^\bullet)^2 \leq L_2 + (\sigma_\zeta^\bullet)^2.$$

At the same time, since $(\sigma_\zeta^\bullet)^2 \geq L_1$, this conditional variance is at least L_1 . Among values of $\hat{\sigma}_\zeta^2 \in [0, \bar{M}_{\sigma_\zeta}^2]$, there exists exactly one such that the conditional variance under

$F_{r^\bullet \frac{\alpha_i + \alpha_{-i}\psi(\kappa^\bullet)}{\alpha_i + \alpha_{-i}\psi(\kappa)}, \kappa, \hat{\sigma}_\zeta}$ is the same as that under $F_{r^\bullet, \kappa^\bullet, \sigma_\zeta^\bullet}$, since we have let $\bar{M}_{\sigma_\zeta}^2 \geq (\sigma_\zeta^\bullet)^2 + L_2$. Thus there is one choice of $\hat{\sigma}_\zeta \in [0, \bar{M}_{\sigma_\zeta}]$ with such that $D_{KL}(F_{r^\bullet, \kappa^\bullet, \sigma_\zeta^\bullet}(\alpha_i, \alpha_{-i}) \parallel F_{r^\bullet \frac{\alpha_i + \alpha_{-i}\psi(\kappa^\bullet)}{\alpha_i + \alpha_{-i}\psi(\kappa)}, \kappa, \hat{\sigma}_\zeta}(\alpha_i, \alpha_{-i})) = 0$. For any other choice of $\tilde{\sigma}_\zeta$, we conclude that $D_{KL}(F_{r^\bullet, \kappa^\bullet, \sigma_\zeta^\bullet}(\alpha_i, \alpha_{-i}) \parallel F_{r^\bullet \frac{\alpha_i + \alpha_{-i}\psi(\kappa^\bullet)}{\alpha_i + \alpha_{-i}\psi(\kappa)}, \kappa, \tilde{\sigma}_\zeta}(\alpha_i, \alpha_{-i})) > 0$. \square

A.10 Proof of Lemma 4

Proof. Assumption A.1 holds as $\mathbb{A}, \Theta_A, \Theta_B$ are compact due to the finite bounds $\bar{M}_\alpha, \bar{M}_r, \bar{M}_{\sigma_\zeta}$. Also, from Lemma 2, the expected utility from playing α_i against α_{-i} in a model with parameters $(\hat{r}, \kappa, \sigma_\zeta)$ is $\mathbb{E}[s_i^2] \cdot (\alpha_i \gamma - \frac{1}{2} \hat{r} \alpha_i^2 - \frac{1}{2} \hat{r} \psi(\kappa) \alpha_i \alpha_{-i} - \frac{1}{2} \alpha_i^2)$. This is a continuous function in $(\alpha_i, \alpha_{-i}, \hat{r})$ and strictly concave in α_i . Therefore Assumptions A.2 and A.5 are satisfied.

To see the finiteness and continuity of the K functions, first recall that the KL divergence from a true distribution $\mathcal{N}(\mu_1, \sigma_1^2)$ to a different distribution $\mathcal{N}(\mu_2, \sigma_2^2)$ is given by $\ln(\sigma_2/\sigma_1) + \frac{\sigma_1^2 + (\mu_1 - \mu_2)^2}{2\sigma_2^2} - \frac{1}{2}$. Under own play α_i , opponent play α_{-i} , correlation parameter κ , elasticity \hat{r} and price idiosyncratic variance σ_ζ^2 , the expected distribution of price after signal s_i is

$$-\frac{1}{2} \hat{r} \alpha_i s_i + (\omega - \frac{1}{2} \hat{r} \alpha_{-i} s_{-i} \mid s_i, \kappa) + \hat{\zeta}$$

where the first term is not random, the middle term is the conditional distribution of $\omega - \frac{1}{2} \hat{r} \alpha_{-i} s_{-i}$ given s_i , based on the joint distribution of (ω, s_i, s_{-i}) with correlation parameter κ . The final term is an independent random variable with mean 0, variance σ_ζ^2 . The analogous true distribution of price is

$$-\frac{1}{2} r^\bullet \alpha_i s_i + (\omega - \frac{1}{2} r^\bullet \alpha_{-i} s_{-i} \mid s_i, \kappa^\bullet) + \zeta^\bullet$$

where ζ^\bullet is an independent random variable with mean 0, variance $(\sigma_\zeta^\bullet)^2$. For a fixed κ , we may find $0 < \underline{\sigma}^2 < \bar{\sigma}^2 < \infty$ so that the variances of both distributions lie in $[\underline{\sigma}^2, \bar{\sigma}^2]$ for all $s_i \in \mathbb{R}$, $\alpha_i, \alpha_{-i} \in [0, \bar{M}_\alpha]$, $\hat{r} \in [0, \bar{M}_r]$. First note that as a consequence of the multivariate normality, the variances of these two expressions do not change with the realization of s_i . The lower bound comes from the fact that $\text{Var}_\kappa(\omega - \frac{1}{2} \hat{r} \alpha_{-i} s_{-i} \mid s_i)$ is nonzero for all α_{-i}, \hat{r} in the compact domains and it is a continuous function of these two arguments, so it must have some positive lower bound $\underline{\sigma}^2 > 0$. For a similar reason, the variance of the middle term has an upper bound for choices of the parameters α_{-i}, \hat{r} in the compact domains, and the inference about σ_ζ^2 is also bounded.

The difference in the means of the two distributions is no larger than $s_i \cdot [\frac{1}{2}(\bar{M}_r + r^\bullet) \cdot 1 + \frac{1}{2}(\bar{M}_r + r^\bullet) \cdot 1 \cdot (\psi(\kappa) + \psi(\kappa^\bullet))]$. Thus consider the function

$$h(s_i) := \ln(\bar{\sigma}/\underline{\sigma}) + \frac{1}{2}(\bar{\sigma}^2/\underline{\sigma}^2) + \frac{[\frac{1}{2}(\bar{M}_r + r^\bullet) \cdot 1 + \frac{1}{2}(\bar{M}_r + r^\bullet) \cdot 1 \cdot (\psi(\kappa) + \psi(\kappa^\bullet))]^2}{2\sigma^2} s_i^2 - \frac{1}{2}.$$

That is $h(s_i)$ has the form $h(s_i) = C_1 + C_2 s_i^2$ for constants C_1, C_2 . It is absolutely integrable against the distribution of s_i , and it dominates the KL divergence between the true and expected price distributions at every s_i and for any choices of $\alpha_i, \alpha_{-i} \in [0, \bar{M}_\alpha]$, $\hat{r} \in [0, \bar{M}_r]$, $\sigma_\zeta^2 \in$

$[0, \bar{M}_\zeta]$. This shows K_A, K_B are finite, so Assumption A.3 holds. Further, since the KL divergence is a continuous function of the means and variances of the price distributions, and since these mean and variance parameters are continuous functions of $\alpha_i, \alpha_{-i}, \hat{r}, \sigma_\zeta^2$, the existence of the absolutely integrable dominating function h also proves K_A, K_B (as integrals of KL divergences across different s_i) are continuous, so Assumption A.4 holds. \square

A.11 Proof of Proposition 5

Proof. We can take L_1, L_2, L_3 as given by Lemma 3. Suppose there is an EZ with behavior $\alpha = (\alpha_{AA}, \alpha_{AB}, \alpha_{BA}, \alpha_{BB})$ and beliefs over parameters $\mu_A \in \Delta(\Theta(\kappa^\bullet)), \mu_B \in \Delta(\Theta(\kappa))$. By Lemma 3, both μ_A and μ_B must be degenerate beliefs that induce zero KL divergence, since both groups match up with group A with probability 1. Furthermore, since Θ_A is correctly specified, it is easy to see that the model $F_{r^\bullet, \kappa^\bullet, \sigma_\zeta^\bullet}$ generates 0 KL divergence, hence the belief of the adherents of Θ_A must be degenerate on this correct model.

In terms of behavior, from Lemma 2, $\alpha_i^{BR}(\alpha_{-i}; \kappa, r) \leq \gamma$ for all $\alpha_{-i} \geq 0, \kappa \in [0, 1], r \geq 0$. Since the upper bound $\bar{M}_\alpha \geq \gamma$, the adherents of each theory must be best responding (across all linear strategies in $[0, \infty)$) in all matches, given their beliefs about the environment.

Using the equilibrium belief of group A, we must have $\alpha_{AA} = \alpha_i^{BR}(\alpha_{AA}; \kappa^\bullet, r^\bullet)$, so $\alpha_{AA} = \frac{\gamma - \frac{1}{2} r^\bullet \psi(\kappa^\bullet) \alpha_{AA}}{1 + r^\bullet}$. We find the unique solution $\alpha_{AA} = \frac{\gamma}{1 + r^\bullet + \frac{1}{2} r^\bullet \psi(\kappa^\bullet)}$.

Next we turn to α_{AB}, α_{BA} , and μ_B . We know μ_B puts probability 1 on some r_B . For adherents of groups A and B to best respond to each others' play and for group B's inference to have 0 KL divergence (when paired with an appropriate choice of σ_ζ), we must have $\alpha_{AB} = \frac{\gamma - \frac{1}{2} r^\bullet \psi(\kappa^\bullet) \alpha_{BA}}{1 + r^\bullet}$, $\alpha_{BA} = \frac{\gamma - \frac{1}{2} r_B \psi(\kappa) \alpha_{AB}}{1 + r_B}$, and $r_B = r^\bullet \frac{\alpha_{BA} + \alpha_{AB} \psi(\kappa^\bullet)}{\alpha_{BA} + \alpha_{AB} \psi(\kappa)}$ from Lemma 3. We may rearrange the expression for α_{BA} to say $\alpha_{BA} = \gamma - r_B \alpha_{BA} - \frac{1}{2} r_B \psi(\kappa) \alpha_{AB}$. Substituting the expression of r_B into this expression of α_{BA} , we get

$$\begin{aligned} \alpha_{BA} &= \gamma - r_B \cdot (\alpha_{BA} + \alpha_{AB} \psi(\kappa) - \frac{1}{2} \alpha_{AB} \psi(\kappa)) \\ &= \gamma - \frac{r^\bullet \alpha_{BA} + r^\bullet \alpha_{AB} \psi(\kappa^\bullet)}{\alpha_{BA} + \alpha_{AB} \psi(\kappa)} \cdot (\alpha_{BA} + \alpha_{AB} \psi(\kappa) - \frac{1}{2} \alpha_{AB} \psi(\kappa)) \\ &= \gamma - r^\bullet \alpha_{BA} - r^\bullet \alpha_{AB} \psi(\kappa^\bullet) + \frac{1}{2} \psi(\kappa) \alpha_{AB} \frac{r^\bullet \alpha_{BA} + r^\bullet \alpha_{AB} \psi(\kappa^\bullet)}{\alpha_{BA} + \alpha_{AB} \psi(\kappa)} \end{aligned}$$

Multiply by $\alpha_{BA} + \alpha_{AB} \psi(\kappa)$ on both sides and collect terms by powers of α ,

$$(\alpha_{BA})^2 \cdot [-1 - r^\bullet] + (\alpha_{BA} \alpha_{AB}) \cdot \left[-\psi(\kappa) - \frac{1}{2} r^\bullet \psi(\kappa) - r^\bullet \psi(\kappa^\bullet) \right] - (\alpha_{AB})^2 \cdot \left[\frac{1}{2} r^\bullet \psi(\kappa^\bullet) \psi(\kappa) \right] + \gamma [\alpha_{BA} + \alpha_{AB} \psi(\kappa)] = 0.$$

Consider the following quadratic function in x ,

$$H(x) := x^2[-1 - r^\bullet] + (x \cdot \ell(x)) \cdot \left[-\psi(\kappa) - \frac{1}{2}r^\bullet\psi(\kappa) - r^\bullet\psi(\kappa^\bullet) \right] - (\ell(x))^2 \cdot \left[\frac{1}{2}r^\bullet\psi(\kappa^\bullet)\psi(\kappa) \right] + \gamma[x + \ell(x)\psi(\kappa)] = 0, \quad (1)$$

where $\ell(x) := \frac{\gamma - \frac{1}{2}r^\bullet\psi(\kappa^\bullet)x}{1+r^\bullet}$ is a linear function in x . In an EZ, α_{BA} is a root of $H(x)$ in $[0, \frac{\gamma}{\frac{1}{2}r^\bullet\psi(\kappa^\bullet)}]$. To see why, if we were to have $\alpha_{BA} > \frac{\gamma}{\frac{1}{2}r^\bullet\psi(\kappa^\bullet)}$, then $\alpha_{AB} = 0$. In that case, $r_B = r^\bullet$ and so $\alpha_{BA} = \alpha_i^{BR}(0, ; \kappa^\bullet, r^\bullet) = \frac{\gamma}{1+r^\bullet}$. Yet $\frac{\gamma}{1+r^\bullet} < \frac{\gamma}{\frac{1}{2}r^\bullet\psi(\kappa^\bullet)}$, contradiction. Conversely, for any root x^* of $H(x)$ in $[0, \frac{\gamma}{\frac{1}{2}r^\bullet\psi(\kappa^\bullet)}]$, there is an EZ where $\alpha_{BA} = x^*$, $\alpha_{AB} = \ell(x^*) \in [0, \gamma]$, and $r_B = r^\bullet \frac{\alpha_{BA} + \alpha_{AB}\psi(\kappa^\bullet)}{\alpha_{BA} + \alpha_{AB}\psi(\kappa)}$.

We now show $H(x)$ (i) has a unique root in $[0, \frac{\gamma}{\frac{1}{2}r^\bullet\psi(\kappa^\bullet)}]$ when $\kappa = \kappa^\bullet$; (ii) does not have a root at $x = 0$ or $x = \frac{\gamma}{\frac{1}{2}r^\bullet\psi(\kappa^\bullet)}$, and (iii) the root in the interval is not a double root. Since $H(x)$ is a continuous function of κ , there must exist some $\underline{\kappa}_1 < \kappa^\bullet < \bar{\kappa}_1$ so that it continues to have a unique root in $[0, \frac{\gamma}{\frac{1}{2}r^\bullet\psi(\kappa^\bullet)}]$ for all $\kappa \in [\underline{\kappa}_1, \bar{\kappa}_1] \cap [0, 1]$.

Claim (i) has to do with the fact that if $\kappa = \kappa^\bullet$, then we need $\alpha_{AB} = \frac{\gamma - \frac{1}{2}r^\bullet\psi(\kappa^\bullet)\alpha_{BA}}{1+r^\bullet}$ and $\alpha_{BA} = \frac{\gamma - \frac{1}{2}r^\bullet\psi(\kappa^\bullet)\alpha_{AB}}{1+r^\bullet}$. These are linear best response functions with a slope of $-\frac{1}{2}\frac{r^\bullet}{1+r^\bullet}\psi(\kappa^\bullet)$, which falls in $(-\frac{1}{2}, 0)$. So there can only be one solution to H in that region (even when we allow $\alpha_{AB} \neq \alpha_{BA}$), which is the symmetric equilibrium found before $\alpha_{AB} = \alpha_{BA} = \frac{\gamma}{1+r^\bullet + \frac{1}{2}r^\bullet\psi(\kappa^\bullet)}$.

For Claim (ii), we evaluate $H(0) = -(\frac{\gamma}{1+r^\bullet})^2 \frac{1}{2}r^\bullet\psi(\kappa^\bullet)^2 + \frac{\gamma^2\psi(\kappa^\bullet)}{1+r^\bullet} = \frac{\psi(\kappa^\bullet)\gamma^2}{1+r^\bullet} (1 - \frac{(1/2)r^\bullet\psi(\kappa^\bullet)}{1+r^\bullet}) \neq 0$ because $1 + r^\bullet > (1/2)r^\bullet\psi(\kappa^\bullet)$. Finally, we evaluate $H(\frac{\gamma}{\frac{1}{2}r^\bullet\psi(\kappa^\bullet)}) = (\frac{\gamma}{\frac{1}{2}r^\bullet\psi(\kappa^\bullet)})^2(-1 - r^\bullet) + \gamma\frac{\gamma}{\frac{1}{2}r^\bullet\psi(\kappa^\bullet)} = \frac{\gamma^2}{\frac{1}{2}r^\bullet\psi(\kappa^\bullet)}(1 - \frac{1+r^\bullet}{\frac{1}{2}r^\bullet\psi(\kappa^\bullet)})$. This is once again not 0 because $1 + r^\bullet > (1/2)r^\bullet\psi(\kappa^\bullet)$.

For Claim (iii), we show that $H'(x^*) < 0$ where $x^* = \frac{\gamma}{1+r^\bullet + \frac{1}{2}r^\bullet\psi(\kappa^\bullet)}$. We find that

$$\begin{aligned} H'(x) = & 2x(-1 - r^\bullet) + \left(\frac{\gamma - r^\bullet\psi(\kappa^\bullet)x}{1+r^\bullet} \right) \left(-\psi(\kappa^\bullet) - \frac{1}{2}r^\bullet\psi(\kappa^\bullet) - r^\bullet\psi(\kappa^\bullet) \right) \\ & - 2 \left(\frac{\gamma - \frac{1}{2}r^\bullet\psi(\kappa^\bullet)x}{1+r^\bullet} \right) \left(\frac{-\frac{1}{2}r^\bullet\psi(\kappa^\bullet)}{1+r^\bullet} \right) \left(\frac{1}{2}r^\bullet\psi(\kappa^\bullet)^2 \right) + \gamma - \frac{\frac{1}{2}r^\bullet\psi(\kappa^\bullet)}{1+r^\bullet}\gamma\psi(\kappa^\bullet). \end{aligned}$$

Collecting terms, the coefficient on x is

$$-2 - 2r^\bullet + \frac{\psi(\kappa^\bullet)^2 r^\bullet}{1+r^\bullet} \left(\frac{3}{2}r^\bullet + 1 - \frac{1}{4} \left(\frac{r^\bullet}{1+r^\bullet} \right)^2 \psi(\kappa^\bullet)^2 \right),$$

while the coefficient on the constant is

$$\frac{\gamma\psi(\kappa^\bullet)}{1+r^\bullet} \left(-\frac{3}{2}r^\bullet - 1 + \frac{1}{2} \frac{(r^\bullet)^2\psi(\kappa^\bullet)^2}{1+r^\bullet} - \frac{1}{2}r^\bullet\psi(\kappa^\bullet) \right) + \gamma.$$

Therefore, we may calculate $H'(x^*) \cdot \frac{1}{x^*}(1+r^\bullet)^2$, which has the same sign as $H'(x^*)$, to be:

$$\begin{aligned} & - (1+r^\bullet)^2(2+2r^\bullet) + \psi(\kappa^\bullet)^2 r^\bullet ((1+r^\bullet)(\frac{3}{2}r^\bullet + 1) - \frac{1}{4}(r^\bullet)^2 \psi(\kappa^\bullet)^2) \\ & + (1+r^\bullet + \frac{1}{2}r^\bullet \psi(\kappa^\bullet)) \left[\psi(\kappa^\bullet)((1+r^\bullet)[-\frac{3}{2}r^\bullet - 1 - \frac{1}{2}r^\bullet \psi(\kappa^\bullet)] + \frac{1}{2}(r^\bullet)^2 \psi(\kappa^\bullet)^2) + (1+r^\bullet)^2 \right]. \end{aligned}$$

We have

$$-(1+r^\bullet)^2(2+2r^\bullet) + (1+r^\bullet + \frac{1}{2}r^\bullet \psi(\kappa^\bullet))(1+r^\bullet)^2 \leq (1+r^\bullet)^2(-1 - \frac{1}{2}r^\bullet) < 0,$$

since $0 \leq \psi(\kappa^\bullet) \leq 1$. Also, for the same reason,

$$(1+r^\bullet)[-\frac{1}{2}r^\bullet \psi(\kappa^\bullet)] + \frac{1}{2}(r^\bullet)^2 \psi(\kappa^\bullet)^2 \leq -\frac{1}{2}(r^\bullet)^2 \psi(\kappa^\bullet) + \frac{1}{2}(r^\bullet)^2 \psi(\kappa^\bullet)^2 \leq 0.$$

Finally, $\psi(\kappa^\bullet)^2 r^\bullet (1+r^\bullet)(\frac{3}{2}r^\bullet + 1) + (1+r^\bullet + \frac{1}{2}r^\bullet \psi(\kappa^\bullet))\psi(\kappa^\bullet)(1+r^\bullet)(-\frac{3}{2}r^\bullet - 1)$ is no larger than

$$\begin{aligned} & \psi(\kappa^\bullet)^2 r^\bullet (\frac{3}{2}(r^\bullet)^2 + \frac{5}{2}r^\bullet + 1) + [r^\bullet \psi(\kappa^\bullet) r^\bullet (-\frac{3}{2}r^\bullet)] \\ & + [r^\bullet \psi(\kappa^\bullet) r^\bullet (-1) + 1 \cdot \psi(\kappa^\bullet) r^\bullet (-\frac{3}{2}r^\bullet)] + [r^\bullet \psi(\kappa^\bullet) \cdot 1 \cdot (-1)] \end{aligned}$$

where the negative terms in the first, second, and third square brackets are respectively larger in absolute value than the first, second and third parts in the expansion of the first summand. Therefore, we conclude $H'(x^*) < 0$.

We have shown that for $\kappa \in [\underline{\kappa}_1, \bar{\kappa}_1] \cap [0, 1]$, group B has only one possible belief about elasticity (denoted by $r_B(\kappa)$) in EZ), since there is only one possible outcome in the match between group A and group B. This means α_{BB} is also pinned down, since there is only one solution to $\alpha_{BB} = \alpha_i^{BR}(\alpha_{BB}, ; \kappa, r_B(\kappa))$. So for every $\kappa \in [\underline{\kappa}_1, \bar{\kappa}_1] \cap [0, 1]$, there is a unique EZ, where equilibrium behavior is given as a function of κ by $\alpha(\kappa) = (\alpha_{AA}(\kappa), \alpha_{AB}(\kappa), \alpha_{BA}(\kappa), \alpha_{BB}(\kappa))$.

Recall from Lemma 2 that the objective expected utility from playing α_i against an opponent who plays α_{-i} is $U_i^\bullet(\alpha_i, \alpha_{-i}) = \mathbb{E}[s_i^2] \cdot (\alpha_i \gamma - \frac{1}{2}r^\bullet \alpha_i^2 - \frac{1}{2}r^\bullet \psi(\kappa^\bullet) \alpha_i \alpha_{-i} - \frac{1}{2}\alpha_i^2)$. If $-i$ plays the rational best response, then the objective expected utility of choosing α_i is $\bar{U}_i(\alpha_i) := \mathbb{E}[s_i^2] \cdot (\alpha_i \gamma - \frac{1}{2}r^\bullet \alpha_i^2 - \frac{1}{2}r^\bullet \psi(\kappa^\bullet) \alpha_i \frac{\gamma - \frac{1}{2}r^\bullet \psi(\kappa^\bullet) \alpha_i}{1+r^\bullet} - \frac{1}{2}\alpha_i^2)$. The derivative in α_i is $\bar{U}'_i(\alpha_i) = \gamma - r^\bullet \alpha_i - \frac{1}{2} \frac{r^\bullet}{1+r^\bullet} \gamma \psi(\kappa^\bullet) + \frac{1}{2} \frac{(r^\bullet)^2 \psi(\kappa^\bullet)^2}{1+r^\bullet} \alpha_i - \alpha_i$. We also know that $\alpha_{AA} = \frac{\gamma}{1+r^\bullet + \frac{1}{2}r^\bullet \psi(\kappa^\bullet)}$

satisfies the first-order condition that $\gamma - r^\bullet \alpha_{AA} - \frac{1}{2} r^\bullet \psi(\kappa^\bullet) \alpha_{AA} - \alpha_{AA} = 0$, therefore

$$\begin{aligned} \bar{U}'_i(\alpha_{AA}) &= -\frac{1}{2} \frac{r^\bullet}{1+r^\bullet} \gamma \psi(\kappa^\bullet) + \frac{1}{2} \frac{(r^\bullet)^2 \psi(\kappa^\bullet)^2}{1+r^\bullet} \alpha_{AA} + \frac{1}{2} r^\bullet \psi(\kappa^\bullet) \alpha_{AA} \\ &= \left[\frac{r^\bullet \psi(\kappa^\bullet)}{2} \right] \left(\frac{-\gamma}{1+r^\bullet} + \frac{\alpha_{AA} \psi(\kappa^\bullet) r^\bullet}{1+r^\bullet} + \alpha_{AA} \right). \end{aligned}$$

Making the substitution $\alpha_{AA} = \frac{\gamma}{1+r^\bullet + \frac{1}{2} r^\bullet \psi(\kappa^\bullet)}$,

$$\begin{aligned} \frac{-\gamma}{1+r^\bullet} + \frac{\alpha_{AA} \psi(\kappa^\bullet) r^\bullet}{1+r^\bullet} + \alpha_{AA} &= \frac{-\gamma(1+r^\bullet + \frac{1}{2} \psi(\kappa^\bullet) r^\bullet) + \gamma \psi(\kappa^\bullet) r^\bullet + \gamma(1+r^\bullet)}{(1+r^\bullet)(1+r^\bullet + \frac{1}{2} \psi(\kappa^\bullet) r^\bullet)} \\ &= \frac{\frac{1}{2} \gamma \psi(\kappa^\bullet) r^\bullet}{(1+r^\bullet)(1+r^\bullet + \frac{1}{2} \psi(\kappa^\bullet) r^\bullet)} > 0. \end{aligned}$$

Therefore, if we can show that $\alpha'_{BA}(\kappa^\bullet) > 0$, then there exists some $\underline{\kappa}_1 \leq \underline{\kappa} < \kappa^\bullet < \bar{\kappa} \leq \bar{\kappa}_1$ so that for every $\kappa \in [\underline{\kappa}, \bar{\kappa}] \cap [0, 1]$, $\kappa \neq \kappa^\bullet$ adherents of Θ_B have strictly higher or strictly lower equilibrium fitness in the unique EZ than adherents of Θ_A , depending on the sign of $\kappa - \kappa^\bullet$. Consider again the quadratic function $H(x)$ in Equation (1) and implicitly characterize the unique root x in $[0, \frac{\gamma}{\frac{1}{2} r^\bullet \psi(\kappa^\bullet)}]$ as a function of κ in a neighborhood around κ^\bullet . Denote this root by α^M , let $D := \frac{d\alpha^M}{d\psi(\kappa)}$ and also note $\frac{d\ell(\alpha^M)}{d\psi(\kappa)} = \frac{-r^\bullet}{2(1+r^\bullet)} \psi(\kappa^\bullet) \cdot D$. We have

$$\begin{aligned} &(-1 - r^\bullet) \cdot (2\alpha^M) \cdot D + (\alpha^M \ell(\alpha^M))(-1 - \frac{1}{2} r^\bullet) \\ &+ (\ell(\alpha^M)D + \alpha^M \frac{-r^\bullet}{2(1+r^\bullet)} \psi(\kappa^\bullet) D) \cdot (-\psi(\kappa) - \frac{1}{2} r^\bullet \psi(\kappa) - r^\bullet \psi(\kappa^\bullet)) + (\ell(\alpha^M))^2 \cdot (-\frac{1}{2} r^\bullet \psi(\kappa^\bullet)) \\ &+ (2\ell(\alpha^M) \frac{-r^\bullet}{2(1+r^\bullet)} \psi(\kappa^\bullet) D) \cdot (-\frac{1}{2} r^\bullet \psi(\kappa^\bullet) \psi(\kappa)) + \gamma(D + \ell(\alpha^M) + \psi(\kappa) \frac{-r^\bullet}{2(1+r^\bullet)} \psi(\kappa^\bullet) D) = 0 \end{aligned}$$

Evaluate at $\kappa = \kappa^\bullet$, noting that $\alpha^M(\kappa^\bullet) = \ell(\alpha^M(\kappa^\bullet)) = x^* := \frac{\gamma}{1+r^\bullet + \frac{1}{2} r^\bullet \psi(\kappa^\bullet) r^\bullet}$.

The terms without D are:

$$\begin{aligned} (x^*)^2(-1 - \frac{1}{2} r^\bullet) + (x^*)^2(\frac{1}{2} r^\bullet \psi(\kappa^\bullet)) + \gamma x^* &= x^* \cdot \left[-x^* \cdot \left(1 + r^\bullet + \frac{1}{2} \psi(\kappa^\bullet) r^\bullet - \frac{1}{2} r^\bullet \right) + \gamma \right] \\ &= x^* \cdot \left[-\gamma + \frac{1}{2} x^* r^\bullet + \gamma \right] = \frac{1}{2} r^\bullet (x^*)^2 > 0. \end{aligned}$$

The coefficient in front of D is:

$$(-1 - r^\bullet)(2x^*) + (x^* + x^* \frac{-r^\bullet}{2(1+r^\bullet)} \psi(\kappa^\bullet)) \cdot (-\psi(\kappa^\bullet) - \frac{3}{2} r^\bullet \psi(\kappa^\bullet)) + \frac{1}{2} x^* \frac{(r^\bullet)^2}{(1+r^\bullet)} \psi(\kappa^\bullet)^3 + \gamma + \gamma \psi(\kappa^\bullet)^2 \cdot \frac{-r^\bullet}{2(1+r^\bullet)}.$$

Make the substitution $\gamma = x^* \cdot \left(1 + r^\bullet + \frac{1}{2}\psi(\kappa^\bullet)r^\bullet\right)$,

$$x^* \cdot \left\{ -2 - 2r^\bullet + \left(1 - \frac{r^\bullet}{2(1+r^\bullet)}\psi(\kappa^\bullet)\right) \cdot \psi(\kappa^\bullet)\left(-\frac{3}{2}r^\bullet - 1\right) + \frac{(r^\bullet)^2}{2(1+r^\bullet)}\psi(\kappa^\bullet)^3 \right\} \\ + x^* \cdot \left\{ \left(1 + r^\bullet + \frac{1}{2}\psi(\kappa^\bullet)r^\bullet\right) \cdot \left(1 - \psi(\kappa^\bullet)^2\frac{r^\bullet}{2(1+r^\bullet)}\right) \right\}.$$

Collect terms inside the parenthesis based on powers of $\psi(\kappa^\bullet)$, we get

$$x^* \cdot \left\{ \psi(\kappa^\bullet)^3 \frac{(r^\bullet)^2}{2(1+r^\bullet)} - \frac{\psi(\kappa^\bullet)^2 r^\bullet}{2(1+r^\bullet)} \left(-\frac{3}{2}r^\bullet - 1\right) + \psi(\kappa^\bullet) \left(-\frac{3}{2}r^\bullet - 1\right) - 2r^\bullet - 2 \right\} \\ + x^* \cdot \left\{ -\psi(\kappa^\bullet)^3 \frac{(r^\bullet)^2}{4(1+r^\bullet)} - \frac{\psi(\kappa^\bullet)^2 r^\bullet}{2(1+r^\bullet)} \cdot (1+r^\bullet) + 1 + r^\bullet + \frac{1}{2}\psi(\kappa^\bullet)r^\bullet \right\}.$$

Combine to get:

$$x^* \cdot \left[\psi(\kappa^\bullet)^3 \frac{(r^\bullet)^2}{4(1+r^\bullet)} + \frac{\psi(\kappa^\bullet)^2 (r^\bullet)^2}{4(1+r^\bullet)} - \psi(\kappa^\bullet)r^\bullet - \psi(\kappa^\bullet) - r^\bullet - 1 \right].$$

Here $\psi(\kappa^\bullet)^3 \frac{(r^\bullet)^2}{4(1+r^\bullet)}$ and $\frac{\psi(\kappa^\bullet)^2 (r^\bullet)^2}{4(1+r^\bullet)}$ are positive terms with

$$\psi(\kappa^\bullet)^3 \frac{(r^\bullet)^2}{4(1+r^\bullet)} + \frac{\psi(\kappa^\bullet)^2 (r^\bullet)^2}{4(1+r^\bullet)} \leq \frac{(r^\bullet)^2}{4(1+r^\bullet)} + \frac{(r^\bullet)^2}{4(1+r^\bullet)} \leq \frac{1}{2} \cdot r^\bullet \cdot \frac{r^\bullet}{1+r^\bullet} \leq \frac{1}{2}r^\bullet.$$

Now $-r^\bullet + \frac{1}{2} \cdot r^\bullet < 0$, and also $-\psi(\kappa^\bullet)r^\bullet - \psi(\kappa^\bullet) - 1 < 0$. Thus the coefficient in front of D is strictly negative. This shows $D(\kappa^\bullet) > 0$. Finally, $\frac{d\alpha^M}{d\psi(\kappa)}$ has the same sign as $\frac{d\alpha^M}{d\kappa}$ since $\psi(\kappa)$ is strictly increasing in κ . \square

A.12 Proof of Proposition 6

Proof. We will show that in every EZ: (i) for each $g \in \{A, B\}$, μ_g puts probability 1 on $\frac{1+\psi(\kappa^\bullet)}{1+\psi(\kappa_g)}r^\bullet$; (ii) for each $g \in \{A, B\}$, $\alpha_{gg} = \frac{\gamma}{1+\frac{r^\bullet}{2}(1+\psi(\kappa^\bullet))+\frac{r^\bullet}{2}\left(\frac{1+\psi(\kappa^\bullet)}{1+\psi(\kappa_g)}\right)}$; (iii) the equilibrium fitness of group A is weakly higher than that of group B if and only if $\kappa_A \leq \kappa_B$.

Choose L_1, L_2, L_3 as in Lemma 3, given r^\bullet and \bar{M}_α . In any EZ with behavior $(\alpha_{AA}, \alpha_{AB}, \alpha_{BA}, \alpha_{BB})$, since the adherents of each theory matches with their own group with probability 1 under perfectly assortatively matching, we conclude that each of μ_g for $g \in \{A, B\}$ must put full weight on $r_i^{INF}(\alpha_{gg}, \alpha_{gg}; \kappa^\bullet, \kappa_g, r^\bullet) = \frac{\alpha_{gg} + \alpha_{gg}\psi(\kappa^\bullet)}{\alpha_{gg} + \alpha_{gg}\psi(\kappa_g)}r^\bullet = \frac{1+\psi(\kappa^\bullet)}{1+\psi(\kappa_g)}r^\bullet$, proving (i).

Given this belief, we must have $\alpha_{gg} = \frac{\gamma - \frac{1}{2}\frac{1+\psi(\kappa^\bullet)}{1+\psi(\kappa_g)}r^\bullet\psi(\kappa_g)\alpha_{gg}}{1 + \frac{1+\psi(\kappa^\bullet)}{1+\psi(\kappa_g)}r^\bullet}$ by Lemma 2. Rearranging yields $\alpha_{gg} = \frac{\gamma}{1 + \frac{r^\bullet}{2}(1+\psi(\kappa^\bullet)) + \frac{r^\bullet}{2}\left(\frac{1+\psi(\kappa^\bullet)}{1+\psi(\kappa_g)}\right)}$, proving (ii).

From Lemma 2, the objective expected utility of each player when both play the strategy profile α_{symm} is $\mathbb{E}[s_i^2] \cdot \left(\alpha_{symm} \gamma - \frac{1}{2} r^\bullet \alpha_{symm}^2 - \frac{1}{2} r^\bullet \psi(\kappa^\bullet) \alpha_{symm}^2 - \frac{1}{2} \alpha_{symm}^2 \right)$. This is a strictly concave quadratic function in α_{symm} that is 0 at $\alpha_{symm} = 0$. Therefore, it is strictly decreasing in α_{symm} for α_{symm} larger than the team solution α_{TEAM} that maximizes this expression, given by the first-order condition

$$\gamma - r^\bullet \alpha_{TEAM} - r^\bullet \psi(\kappa^\bullet) \alpha_{TEAM} - \alpha_{TEAM} = 0 \Rightarrow \alpha_{TEAM} = \frac{\gamma}{1 + r^\bullet + r^\bullet \psi(\kappa^\bullet)}.$$

For any value of $\kappa \in [0, 1]$, using the fact that $\psi(0) > 0$ and ψ is strictly increasing,

$$\frac{\gamma}{1 + \frac{r^\bullet}{2}(1 + \psi(\kappa^\bullet)) + \frac{r^\bullet}{2} \left(\frac{1 + \psi(\kappa^\bullet)}{1 + \psi(\kappa)} \right)} > \frac{\gamma}{1 + \frac{r^\bullet}{2}(1 + \psi(\kappa^\bullet)) + \frac{r^\bullet}{2}(1 + \psi(\kappa^\bullet))} = \alpha_{TEAM}.$$

Also, $\frac{\gamma}{1 + \frac{r^\bullet}{2}(1 + \psi(\kappa^\bullet)) + \frac{r^\bullet}{2} \left(\frac{1 + \psi(\kappa^\bullet)}{1 + \psi(\kappa)} \right)}$ is a strictly increasing function in κ , since ψ is strictly increasing. We therefore conclude that each player's utility when they play $\frac{\gamma}{1 + \frac{r^\bullet}{2}(1 + \psi(\kappa^\bullet)) + \frac{r^\bullet}{2} \left(\frac{1 + \psi(\kappa^\bullet)}{1 + \psi(\kappa)} \right)}$ against each other is strictly decreasing in κ , proving (iii). \square

A.13 Proof of Proposition 7

Proof. Find L_1, L_2, L_3 as given by Lemma 3. Suppose $\Theta_A = \Theta(\kappa^\bullet)$, $\Theta_B = \{F_{r^\bullet, \kappa, \sigma_\zeta^\bullet}\}$ for any $\kappa \in [0, 1]$, $(p_A, p_B) = (1, 0)$, and $\lambda \in [0, 1]$, then arguments similar to those in the proof of Lemma 3 imply there exists exactly one EZ, and it involves the adherents of Θ_A holding correct beliefs and playing $\frac{\gamma}{1 + r^\bullet + \frac{1}{2} r^\bullet \psi(\kappa^\bullet)}$ against each other.

We now analyze $\alpha_{BA}(\kappa)$ in such EZ. In the proof of Proposition 5, we defined $\bar{U}_i(\alpha_i)$ as i 's objective expected utility of choosing α_i when $-i$ plays the rational best response. We showed that $\bar{U}_i' \left(\frac{\gamma}{1 + r^\bullet + \frac{1}{2} r^\bullet \psi(\kappa^\bullet)} \right) > 0$. In an EZ where i believes in the model $F_{r^\bullet, \kappa, \sigma_\zeta^\bullet}$ and $-i$ believes in the model $F_{r^\bullet, \kappa, \sigma_\zeta^\bullet}$, using the expression for α_i^{BR} from Lemma 2, the play of i solves $x = \frac{\gamma - \frac{1}{2} r^\bullet \psi(\kappa) \left(\frac{\gamma - \frac{1}{2} r^\bullet \psi(\kappa^\bullet) x}{1 + r^\bullet} \right)}{1 + r^\bullet}$, which implies $\alpha_{BA}(\kappa) = \frac{\gamma(1 + r^\bullet - \frac{1}{2} \psi(\kappa) r^\bullet)}{1 + 2r^\bullet + (r^\bullet)^2 - \frac{1}{4} \psi(\kappa) \psi(\kappa^\bullet) (r^\bullet)^2}$. Taking the derivative and evaluating at $\kappa = \kappa^\bullet$, we find an expression with the same sign as $\frac{1}{4} \psi'(\kappa^\bullet) r^\bullet (1 + r^\bullet) \gamma (-2(1 + r^\bullet) + \psi(\kappa^\bullet) r^\bullet)$, which is strictly negative because $\psi'(\kappa^\bullet) > 0$, $r^\bullet > 0$, $\gamma > 0$, and $\psi(\kappa^\bullet) \leq 1$. This shows there exists $\epsilon > 0$ so that for every $\kappa_h \in (\kappa^\bullet, \kappa^\bullet + \epsilon]$, we have $\bar{U}_i(\alpha_{BA}(\kappa_h)) < \bar{U}_i \left(\frac{\gamma}{1 + r^\bullet + \frac{1}{2} r^\bullet \psi(\kappa^\bullet)} \right)$, that is the adherents of $\{F_{r^\bullet, \kappa_h, \sigma_\zeta^\bullet}\}$ have strictly lower fitness than the adherents of $\Theta(\kappa^\bullet)$ with $\lambda = 0$ in the unique EZ. Finally, existence and upper-hemicontinuity of EZ in population proportion in such societies can be established using arguments similar to the proof of Propositions A.1 and A.2. This establishes the first claim to be proved.

Next, we turn to $\alpha_{BB}(\kappa)$. Using the expressing for α_i^{BR} in Lemma 2, we find that $\alpha_{BB}(\kappa) =$

$\frac{\gamma}{1+r^\bullet+\frac{1}{2}r^\bullet\psi(\kappa)}$. Since $\psi' > 0$, we have $\alpha_{BB}(\kappa)$ is strictly larger than $\alpha_{AA} = \frac{\gamma}{1+r^\bullet+\frac{1}{2}r^\bullet\psi(\kappa^\bullet)}$ when $\kappa < \kappa^\bullet$. From the proof of Proposition 6, we know that objective payoffs in the stage game is strictly decreasing in linear strategies larger than the team solution $\alpha_{TEAM} = \frac{\gamma}{1+r^\bullet+r^\bullet\psi(\kappa^\bullet)}$. Since $\alpha_{BB}(\kappa) > \alpha_{AA} > \alpha_{TEAM}$, we conclude the adherents of $\{F_{r^\bullet, \kappa_l, \sigma_\zeta^\bullet}\}$ have strictly lower fitness than the adherents of $\Theta(\kappa^\bullet)$ with $\lambda = 1$ in the unique EZ, for any $\kappa_l < \kappa^\bullet$. Again, existence and upper-hemicontinuity of EZ in population proportion in such societies can be established using arguments similar to the proof of Propositions A.1 and A.2. This establishes the second claim to be proved. \square

A.14 Proof of Proposition 8

Proof. Consider the society where $\Theta_A = \Theta_B = \Theta(\kappa^\bullet)$, $(p_A, p_B) = (1, 0)$. For any EZ with behavior $(\sigma_{AA}, \sigma_{AB}, \sigma_{BA}, \sigma_{BB})$ and beliefs (μ_A, μ_B) , there exists another EZ $(\sigma'_{AA}, \sigma'_{AB}, \sigma'_{BA}, \sigma'_{BB})$ where $\sigma'_{g,g'} = \sigma_{AA}$ for all $g, g' \in \{A, B\}$ and all agents hold the belief μ_A . The uniqueness of EZ from Assumption 1 implies $\alpha_{AB}(\kappa^\bullet) = \alpha_{BA}(\kappa^\bullet) = \alpha_{BB}(\kappa^\bullet) = \alpha^\bullet$.

Now consider the society where $\Theta_B = \Theta(\kappa)$, $(p_A, p_B) = (1, 0)$. By the same arguments as the existence arguments in Proposition A.1, there exists an EZ where $\alpha_{AA}(\kappa) = \alpha_{AA}(\kappa^\bullet)$. By the uniqueness of EZ from Assumption 1, we must in fact have $\alpha_{AA}(\kappa) = \alpha_{AA}(\kappa^\bullet)$ for all κ , so the fitness of theory $\Theta(\kappa^\bullet)$ in the unique EZ is

$$\mathbb{E}^\bullet [\mathbb{E}^\bullet [u_1^\bullet(\alpha^\bullet s_1, \alpha^\bullet s_2, \omega) \mid s_1]].$$

Under λ matching with mutant theory $\Theta(\kappa)$, the mutant's fitness in the unique EZ is

$$\mathbb{E}^\bullet [\mathbb{E}^\bullet [(1 - \lambda)u_1^\bullet(\alpha_{BA}(\kappa)s_1, \alpha_{AB}(\kappa)s_2, \omega) + (\lambda)u_1^\bullet(\alpha_{BB}(\kappa)s_1, \alpha_{BB}(\kappa)s_2, \omega) \mid s_1]].$$

Differentiate and evaluate at $\kappa = \kappa^\bullet$. At $\kappa = \kappa^\bullet$, adherents of Θ_A and Θ_B have the same fitness since they play the same strategies. So, a non-zero sign on the derivative would give the desired evolutionary fragility against either theories with slightly higher or slightly lower κ . This derivative is:

$$\mathbb{E}^\bullet \left[\mathbb{E}^\bullet \left[\begin{array}{l} \frac{\partial u_1^\bullet}{\partial q_1}(\alpha^\bullet s_1, \alpha^\bullet s_2, \omega) \cdot [(1 - \lambda)\alpha'_{BA}(\kappa^\bullet) + \lambda\alpha'_{BB}(\kappa^\bullet)] \cdot s_1 \\ + \frac{\partial u_1^\bullet}{\partial q_2}(\alpha^\bullet s_1, \alpha^\bullet s_2, \omega) \cdot [(1 - \lambda)\alpha'_{AB}(\kappa^\bullet) + \lambda\alpha'_{BB}(\kappa^\bullet)] \cdot s_2 \end{array} \middle| s_1 \right] \right].$$

Using the interim optimality part of Assumption 1, $\mathbb{E}^\bullet \left[\frac{\partial u_1^\bullet}{\partial q_1}(\alpha^\bullet s_1, \alpha^\bullet s_2, \omega) \mid s_1 \right] = 0$ for every $s_1 \in S$, using the necessity of the first-order condition. The derivative thus simplifies as claimed. \square

A.15 Proof of Proposition 10

Proof. In the centipede game, suppose $g > \frac{2}{K-2}\ell$. the misspecified agent thinks a group B agent in the role of P2 and a group A agent in either role has a probability $2/K$ of stopping at every node. Under this belief, choosing to continue instead of drop means there is a $(K-2)/K$ chance of gaining g , but a $2/K$ chance of losing ℓ . Since we assume $g > \frac{2}{K-2}\ell$, it is strictly better to continue. When p fraction of the agents are correctly specified, the fitness of Θ^\bullet is $p \cdot 0 + (1-p) \cdot (\frac{1}{2}\frac{g(K-2)}{2} + \frac{1}{2}(\frac{gK}{2} + \ell))$, while the fitness of Θ^{An} is $p \cdot [\frac{1}{2}(\frac{g(K-2)}{2} - \ell) + \frac{1}{2}\frac{g(K-2)}{2}] + (1-p)[\frac{1}{2}(\frac{g(K-2)}{2} - \ell) + \frac{1}{2}(\frac{gK}{2} + \ell)]$. The difference in fitness is

$$-p[\frac{1}{2}(\frac{g(K-2)}{2} - \ell) + \frac{1}{2}\frac{g(K-2)}{2}] + (1-p)\frac{1}{2}\ell.$$

Simplifying, this is $\frac{1}{2}\ell - p \cdot \frac{g(K-2)}{2}$, a strictly decreasing function in p . When $p = \frac{\ell}{g(K-2)}$, which is a number strictly between 0 and $1/2$ from the assumption $g > \frac{2}{K-2}\ell$ in the centipede game, the two theories have the same fitness. \square

A.16 Proof of Proposition 11

Proof. In the $\bar{\Theta}^{An}$ vs. $\bar{\Theta}^{An}$ match, the adherents of $\bar{\Theta}^{An}$ hold the belief that $\hat{d}_{BB}^k = 2/K$ for every even k . In the role of P1, at node k for $k \leq K-3$, stopping gives them k but continuing gives them a $(K-2)/K$ chance to get at least $k+2$, and we have $k \leq \frac{K-2}{K}(k+2) \iff 2k \leq 2K-4 \iff k \leq K-2$. At node $K-1$, the agent gets $K-1$ from dropping but expects $(K+2) \cdot \frac{K-2}{K}$ from continuing, and $(K+2) \cdot \frac{K-2}{K} - (K-1) = \frac{K^2-4-K^2+K}{K} = \frac{K-4}{K} > 0$ since $K \geq 6$.

In the $\bar{\Theta}^\bullet$ vs. $\bar{\Theta}^{An}$ match, the adherents of $\bar{\Theta}^{An}$ hold the belief that $\hat{d}_{AB}^k = 2/K$ for every k . By the same arguments as before, the behavior of the adherents of $\bar{\Theta}^{An}$ are optimal given these beliefs. Also, the adherents of $\bar{\Theta}^\bullet$ have no profitable deviations since they are best responding both as P1 and P2.

When p fraction of the agents are correctly specified, in the dollar game the fitness of $\bar{\Theta}^\bullet$ is $p \cdot 0.5 + (1-p) \cdot (\frac{1}{2}(K-1) + \frac{1}{2}K)$, while the fitness of $\bar{\Theta}^{An}$ is $p \cdot 0 + (1-p) \cdot (\frac{1}{2} \cdot 0 + \frac{1}{2}K)$. For any p , the fitness of $\bar{\Theta}^\bullet$ is strictly higher than that of $\bar{\Theta}^{An}$. \square

B Existence and Continuity of EZ

We provide a few technical results about the existence of EZ and the upper-hemicontinuity of the set of EZ with respect to population share. The existence and continuity results also establish the existence of approachable EZs with population shares $p = (1, 0)$. Note that

the same learning channel that generates new stability phenomena in Section 3 also leads to some difficulty in establishing existence and continuity results, as agents draw different inferences with different interaction structures.

Let two theories, Θ_A, Θ_B be fixed. Also fix population shares p and matching assortativity λ . Let $U_A : \mathbb{A}^2 \times \Theta_A \rightarrow \mathbb{R}$ be such that $U_A(a_i, a_{-i}; F) = U_i(a_i, a_{-i}; \delta_F)$ and let $U_B : \mathbb{A}^2 \times \Theta_B \rightarrow \mathbb{R}$ be such that $U_B(a_i, a_{-i}; F) = U_i(a_i, a_{-i}; \delta_F)$.

Assumption A.1. $\mathbb{A}, \Theta_A, \Theta_B$ are compact metrizable spaces.

Assumption A.2. U_A, U_B are continuous.

Assumption A.3. For every $F \in \Theta_A \cup \Theta_B$ and $a_i, a_{-i} \in \mathbb{A}$, $K(F; a_i, a_{-i})$ is well-defined and finite.

Under Assumption A.3, we have the well-defined functions $K_A : \Theta_A \times \mathbb{A}^2 \rightarrow \mathbb{R}_+$ and $K_B : \Theta_B \times \mathbb{A}^2 \rightarrow \mathbb{R}_+$, where $K_g(F; a_i, a_{-i}) := KL(F^\bullet(a_i, a_{-i}) \parallel F(a_i, a_{-i}))$.

Assumption A.4. K_A and K_B are continuous.

Assumption A.5. \mathbb{A} is convex and, for all $a_{-i} \in \mathbb{A}$ and $\mu \in \Delta(\Theta_A) \cup \Delta(\Theta_B)$, $a_i \mapsto U_i(a_i, a_{-i}; \mu)$ is quasiconcave.

We show existence of EZ using the Kakutani-Fan-Glicksberg fixed point theorem, applied to the correspondence which maps strategy profiles and beliefs over models into best replies and beliefs over KL-divergence minimizing models. We start with a lemma.

Lemma A.1. For $g \in \{A, B\}$, $a = (a_{AA}, a_{AB}, a_{BA}, a_{BB}) \in \mathbb{A}^4$, and $0 \leq m_g \leq 1$, let

$$\Theta_g^*(a, m_g) := \arg \min_{\hat{F} \in \Theta_g} \left\{ m_g \cdot K(\hat{F}; a_{g,g}, a_{g,g}) + (1 - m_g) \cdot K(\hat{F}; a_{g,-g}, a_{-g,g}) \right\}.$$

Then, Θ_g^* is upper hemicontinuous in its arguments.

This lemma says the set of KL-minimizing models is upper hemicontinuous in strategy profile and matching assortativity. This leads to the existence result.

Proposition A.1. Under Assumptions A.1, A.2, A.3, A.4, and A.5, an EZ exists.

Next, upper hemicontinuity in m_g in Lemma A.1 allows us to deduce the upper hemicontinuity of the EZ correspondence in population shares, and conclude that the notion of approachability from Definition 3 is a non-empty refinement of the set of EZ with $p = (1, 0)$.

Proposition A.2. Fix two theories Θ_A, Θ_B . Also fix matching assortativity $\lambda \in [0, 1]$. The set of EZ is an upper hemicontinuous correspondence in p_B under Assumptions A.1, A.2, A.3, and A.4.

Corollary A.1. Under Assumptions A.1, A.2, A.3, A.4, and A.5, the set of approachable EZ with $p = (1, 0)$ is non-empty for every λ .

B.1 Proofs of Results in Appendix B

B.1.1 Proof of Lemma A.1

Proof. Write the minimization objective as

$$W(a, F, m_g) := m_g K_g(F; a_{g,g}, a_{g,g}) + (1 - m_g) K_g(F; a_{g,-g}, a_{-g,g}),$$

a continuous function of (a, F, m_g) by Assumption A.4. Suppose we have a sequence $(a^{(n)}, m_g^{(n)}) \rightarrow (a^*, m_g^*) \in \mathbb{A}^4 \times [0, 1]$ and let $F^{(n)} \in \Theta_g^*(a^{(n)}, m_g^{(n)})$ for each n , with $F^{(n)} \rightarrow F^* \in \Theta_g$. For any other $\hat{F} \in \Theta_g$, note that $W(a^*, m_g^*, \hat{F}) = \lim_{n \rightarrow \infty} W(a^{(n)}, m_g^{(n)}, \hat{F})$ by continuity. But also by continuity, $W(a^*, m_g^*, F^*) = \lim_{n \rightarrow \infty} W(a^{(n)}, m_g^{(n)}, F^{(n)})$ and $W(a^{(n)}, m_g^{(n)}, F^{(n)}) \leq W(a^{(n)}, m_g^{(n)}, \hat{F})$ for every n . It therefore follows $W(a^*, m_g^*, F^*) \leq W(a^*, m_g^*, \hat{F})$. \square

B.1.2 Proof of Proposition A.1

Proof. Consider the correspondence $\Gamma : \mathbb{A}^4 \times \Delta(\Theta_A) \times \Delta(\Theta_B) \rightrightarrows \mathbb{A}^4 \times \Delta(\Theta_A) \times \Delta(\Theta_B)$,

$$\begin{aligned} \Gamma(a_{AA}, a_{AB}, a_{BA}, a_{BB}, \mu_A, \mu_B) := \\ (\text{BR}(a_{AA}, \mu_A), \text{BR}(a_{BA}, \mu_A), \text{BR}(a_{AB}, \mu_B), \text{BR}(a_{BB}, \mu_B), \Delta(\Theta_A^*(a)), \Delta(\Theta_B^*(a))), \end{aligned}$$

where $\text{BR}(a_{-i}, \mu_g) := \arg \max_{\hat{a}_i \in \mathbb{A}} U_g(\hat{a}_i, a_{-i}; \mu_g)$ and, for each $g \in \{A, B\}$, the correspondence Θ_g^* is defined with $m_g = \lambda + (1 - \lambda)p_g$, $m_{-g} = 1 - m_g$. It is clear that fixed points of Γ are EZ.

We apply the Kakutani-Fan-Glicksberg theorem (see, e.g, Corollary 17.55 in [Aliprantis and Border \(2006\)](#)). By Assumptions A.1 and A.5, \mathbb{A} is a compact and convex metric space, and each Θ_g is a compact metric space, so it follows the domain of Γ is a nonempty, compact and convex metric space. We need only verify that Γ has closed graph, non-empty values, and convex values.

To see that Γ has closed graph, the previous lemma shows the upper hemicontinuity of $\Theta_A^*(a)$ and $\Theta_B^*(a)$ in a , and Theorem 17.13 of [Aliprantis and Border \(2006\)](#) then implies $\Delta(\Theta_A^*(a))$ and $\Delta(\Theta_B^*(a))$ are also upper hemicontinuous in a . It is a standard argument that since Assumption A.2 supposes U_A, U_B are continuous, it implies the best-response correspondences $\text{BR}(a_{AA}, \mu_A)$, $\text{BR}(a_{BA}, \mu_A)$, $\text{BR}(a_{AB}, \mu_B)$, $\text{BR}(a_{BB}, \mu_B)$ have closed graphs.

To see that Γ is non-empty, recall that each $\hat{a}_i \mapsto U_g(\hat{a}_i, a_{-i}; \mu_g)$ is a continuous function on a compact domain, so it must attain a maximum on \mathbb{A} . Similarly, the minimization problem that defines each $\Theta_g^*(a)$ is a continuous function of F over a compact domain of possible F 's, so it attains a minimum. Thus each $\Delta(\Theta_g^*(a))$ is the set of distributions over a non-empty set.

To see that Γ is convex valued, clearly $\Delta(\Theta_A^*(a))$ and $\Delta(\Theta_B^*(a))$ are convex valued by definition. Also, $\hat{a}_i \mapsto U_A(\hat{a}_i, a_{AA}; \mu_A)$ is quasiconcave by Assumption A.5. That means if $a'_i, a''_i \in \text{BR}(a_{AA}, \mu_A)$, then for any convex combination \tilde{a}_i of a'_i, a''_i , we have $U_A(\tilde{a}_i, a_{AA}; \mu_A) \geq \min(U_A(a'_i, a_{AA}; \mu_A), U_A(a''_i, a_{AA}; \mu_A)) = \max_{\hat{a}_i \in \mathbb{A}} U_A(\hat{a}_i, a_{AA}; \mu_A)$. Therefore, $\text{BR}(a_{AA}, \mu_A)$ is convex. For similar reasons, $\text{BR}(a_{BA}, \mu_A)$, $\text{BR}(a_{AB}, \mu_B)$, $\text{BR}(a_{BB}, \mu_B)$ are convex. \square

B.1.3 Proof of Proposition A.2

Proof. Since $\mathbb{A}^4 \times \Delta(\Theta_A) \times \Delta(\Theta_B)$ is compact by Assumption A.1, we need only show that for every sequence $(p_B^{(k)})_{k \geq 1}$ and $(a^{(k)}, \mu^{(k)})_{k \geq 1} = (a_{AA}^{(k)}, a_{AB}^{(k)}, a_{BA}^{(k)}, a_{BB}^{(k)}, \mu_A^{(k)}, \mu_B^{(k)})_{k \geq 1}$ such that for every k , $(a^{(k)}, \mu^{(k)})$ is an EZ with $p = (1 - p_B^{(k)}, p_B^{(k)})$, $p_B^{(k)} \rightarrow p_B^*$, and $(a^{(k)}, \mu^{(k)}) \rightarrow (a^*, \mu^*)$, then (a^*, μ^*) is an EZ with $p = (1 - p_B^*, p_B^*)$.

We first show for all $g, g' \in \{A, B\}$, $a_{g,g'}^*$ is optimal against $a_{g',g}^*$ under the belief μ_g^* . Assortativity does not matter here, since optimality applies within all type match-ups. By Assumption A.2, $U_g(a_i, a_{-i}; F)$ is continuous, so by property of convergence in distribution, $U_g(a_{g,g'}^{(k)}, a_{g',g}^{(k)}; \mu_g^{(k)}) \rightarrow U_g(a_{g,g'}^*, a_{g',g}^*; \mu_g^*)$. For any other $\hat{a}_i \in \mathbb{A}$, $U_g(\hat{a}_i, a_{g',g}^{(k)}; \mu_g^{(k)}) \rightarrow U_g(\hat{a}_i, a_{g',g}^*; \mu_g^*)$ and for every k , $U_g(a_{g,g'}^{(k)}, a_{g',g}^{(k)}; \mu_g^{(k)}) \geq U_g(\hat{a}_i, a_{g',g}^{(k)}; \mu_g^{(k)})$. Therefore $a_{g,g'}^*$ best responds to $a_{g',g}^*$ under belief μ_g^* .

Next, we show models in the support of μ_g^* minimize weighted KL divergence for group g . First consider the correspondence $H : \mathbb{A}^4 \times [0, 1] \rightrightarrows \Theta_g$ where $H(a, p_g) := \Theta_g^*(a, \lambda + (1 - \lambda)(p_g))$. Then H is upper hemicontinuous by Lemma A.1. Since $H(a, p_g)$ represents the minimizers of a continuous function on a compact domain, it is non-empty and closed. By Theorem 17.13 of Aliprantis and Border (2006), the correspondence $\tilde{H} : \mathbb{A}^4 \times [0, 1] \rightrightarrows \Delta(\Theta_g)$ defined so that $\tilde{H}(a, p_g) := \Delta(H(a, p_g))$ is also upper hemicontinuous. For every k , $\mu_g^{(k)} \in \tilde{H}(a^{(k)}, p_g^{(k)})$, and $\mu_g^{(k)} \rightarrow \mu_g^*$, $a^{(k)} \rightarrow a^*$, $p_g^{(k)} \rightarrow p_g^*$. Therefore, $\mu_g^* \in \tilde{H}(a^*, p_g^*)$, that is to say μ_g^* is supported on the minimizers of weighted KL divergence. \square

C Learning Foundation of EZ and EZ-SU

We provide a foundation for EZ and EZ-SU as the steady state of a learning system.

We first consider a world where agents have a prior over extended models and do not observe ex-post signals about the matched opponent's strategy. We show that the steady states in this world correspond to EZ-SU. Then, we show that the addition of sufficiently accurate ex-post signals about opponent's strategy implies the steady states are EZ.

C.1 Regularity Assumptions

We make some regularity assumptions on the objective environments and on the theories Θ_A, Θ_B . These are similar to the regularity assumptions from Section B.

Suppose \mathbb{A} is finite. Suppose the marginals of $\bar{\Theta}_A, \bar{\Theta}_B$ on the dimension of fundamental uncertainty, Θ_A, Θ_B , are compact metrizable spaces. Endow $\bar{\Theta}_A$ and $\bar{\Theta}_B$ with the product metric. Suppose that every $(a_A, a_B, F) \in \bar{\Theta}_A \cup \bar{\Theta}_B$ is so that for every $(a_i, a_{-i}) \in \mathbb{A}^2$, whenever $f^\bullet(a_i, a_{-i})(y) > 0$, we also get $f(a_i, a_A)(y) > 0$ and $f(a_i, a_B)(y) > 0$, where f is the density or probability mass function for F .

For each $g, g' \in \{A, B\}$, define $K_{g,g'} : \mathbb{A}^2 \times \bar{\Theta}_g \rightarrow \mathbb{R}$ by $K_{g,g'}(a_i, a_{-i}; (a_A, a_B, F)) = KL(F^\bullet(a_i, a_{-i}) \parallel F(a_i, a_{g'}))$. Suppose each $K_{g,g'}$ is well defined and a continuous function of the extended model (a_A, a_B, F) .

For $g \in \{A, B\}$, $F \in \Theta_g$, let $U_g(a_i, a_{-i}; F)$ be the expected payoffs of the strategy profile (a_i, a_{-i}) for i when consequences are drawn according to F . Assume U_A, U_B are continuous.

Suppose for every extended theory $\bar{\Theta}_g$ and every $(a_A, a_B, F) \in \bar{\Theta}_g$ and $\epsilon > 0$, there exists an open neighborhood $V \subseteq \bar{\Theta}_g$ of (a_A, a_B, F) , so that for every $(\hat{a}_A, \hat{a}_B, \hat{F}) \in V$, $1 - \epsilon \leq f(a_i, a_A)(y)/\hat{f}(a_i, \hat{a}_A)(y) \leq 1 + \epsilon$ and $1 - \epsilon \leq f(a_i, a_B)(y)/\hat{f}(a_i, \hat{a}_B)(y) \leq 1 + \epsilon$ for all $a_i \in \mathbb{A}, y \in \mathbb{Y}$. Also suppose there is some $M > 0$ so that $\ln(f(a_i, a_A)(y))$ and $\ln(f(a_i, a_B)(y))$ are bounded in $[-M, M]$ for all $(a_A, a_B, F) \in \bar{\Theta}_g, a_i, a_{-i} \in \mathbb{A}, y \in \mathbb{Y}$.

C.2 Learning Environment

Time is discrete and infinite, $t = 0, 1, 2, \dots$. A unit mass of agents, $i \in [0, 1]$, enter the society at time 0. A $p_A \in (0, 1)$ measure of them are assigned to theory A and the rest are assigned to theory B . Each agent born into theory g starts with the same full support prior over the extended theory, $\mu_g^{(0)} \in \Delta(\bar{\Theta}_g)$, and believes there is some $(a_A, a_B, F) \in \bar{\Theta}_g$ so that every group g opponent always plays a_g and the consequences are always generated by F .

In each period t , agents are matched up partially assortatively to play the stage game. Assortativity is $\lambda \in (0, 1)$. Each person in group g has $\lambda + (1 - \lambda)p_g$ chance of matching with someone from group g , and matches with someone from group $-g$ with the complementary chance. Each agent i observes their opponent's group membership and chooses a strategy $a_i^{(t)} \in \mathbb{A}$. At the end of the match, the agent observes own consequence $y_i^{(t)}$ and an ex-post signal $x_i^{(t)} \in \mathbb{A}$, where $x_i^{(t)}$ equals the matched opponent's strategy a_{-i} with probability $\tau \in [0, 1]$, and it is uniformly random on \mathbb{A} with the complementary probability. To give a foundation for a EZ-SU, we consider $\tau = 0$, so the signal x_i is uninformative. To give a foundation for EZ, we consider τ close to 1.

Thus, the space of histories from one period is $\{A, B\} \times \mathbb{A} \times \mathbb{Y} \times \mathbb{A}$, where the first instance of the strategy is own strategy and the second instance is the ex-post signal. Let

\mathbb{H} denote the space of all finite-length histories.

Given the assumption on the two theories, there is a well-defined Bayesian belief operator for each theory g , $\mu_g : \mathbb{H} \rightarrow \Delta(\bar{\Theta}_g)$, mapping every finite-length history into a belief over extended models in $\bar{\Theta}_g$, starting with the prior $\mu_g^{(0)}$.

We also take as exogenously given policy functions for choosing strategies after each history. That is, $\mathbf{a}_{g,g'} : \mathbb{H} \rightarrow \mathbb{A}$ for every $g, g' \in \{A, B\}$ gives the strategy that a group g agent uses against a group g' opponent after every history. Assume these policy functions are asymptotically myopic.

Assumption A.6. *For every $\epsilon > 0$, there exists K so that for any history h containing at least K matches against opponents of each group, $\mathbf{a}_{g,g'}(h)$ is an ϵ -best response to the Bayesian belief $\mu_g(h)$ about the model.*

From the perspective of each agent i in group g , i 's play against groups A and B, as well as i 's belief over $\bar{\Theta}_g$, is a stochastic process $(\tilde{a}_{iA}^{(t)}, \tilde{a}_{iB}^{(t)}, \tilde{\mu}_i^{(t)})_{t \geq 0}$ valued in $\mathbb{A} \times \mathbb{A} \times \Delta(\bar{\Theta}_g)$. The randomness is over the groups of opponents matched with in different periods, the strategies they play, and the random consequence and ex-post signals drawn at the end of the match. At the same time, since there is a continuum of agents, the distribution over histories within each population in each period is deterministic. As such, there is a deterministic sequence $(\alpha_{AA}^{(t)}, \alpha_{AB}^{(t)}, \alpha_{BA}^{(t)}, \alpha_{BB}^{(t)}, \nu_A^{(t)}, \nu_B^{(t)}) \in \Delta(\mathbb{A})^4 \times \Delta(\Delta(\bar{\Theta}_A)) \times \Delta(\Delta(\bar{\Theta}_B))$ that describes the distributions of play and beliefs that prevail in the two sub-populations in every period t .

C.3 Steady State Limits are EZ-SUs and EZs

We state and prove the learning foundation of EZ-SU and EZ. For $(\alpha^{(t)})_t$ a sequence valued in $\Delta(\mathbb{A})$ and $a^* \in \mathbb{A}$, $\alpha^{(t)} \rightarrow a^*$ means $\mathbb{E}_{\hat{a} \sim \alpha^{(t)}} \|\hat{a} - a^*\| \rightarrow 0$ as $t \rightarrow \infty$. For $(\nu^{(t)})_t$ a sequence valued in $\Delta(\Delta(\bar{\Theta}_g))$ and $\mu^* \in \Delta(\bar{\Theta}_g)$, $\nu^{(t)} \rightarrow \mu^*$ means $\mathbb{E}_{\hat{\mu} \sim \nu^{(t)}} \|\hat{\mu} - \mu^*\| \rightarrow 0$ as $t \rightarrow \infty$.

Proposition A.3. *Suppose the regularity assumptions in Section C.1 hold, and suppose Assumption A.6 holds.*

Suppose $\tau = 0$. Suppose there exists $(a_{AA}^, a_{AB}^*, a_{BA}^*, a_{BB}^*, \mu_A^*, \mu_B^*) \in \mathbb{A}^4 \times \Delta(\bar{\Theta}_A) \times \Delta(\bar{\Theta}_B)$ so that $(\alpha_{AA}^{(t)}, \alpha_{AB}^{(t)}, \alpha_{BA}^{(t)}, \alpha_{BB}^{(t)}, \nu_A^{(t)}, \nu_B^{(t)}) \rightarrow (a_{AA}^*, a_{AB}^*, a_{BA}^*, a_{BB}^*, \mu_A^*, \mu_B^*)$ and for each agent i in group g , almost surely $(\tilde{a}_{iA}^{(t)}, \tilde{a}_{iB}^{(t)}, \tilde{\mu}_i^{(t)}) \rightarrow (a_{gA}^*, a_{gB}^*, \mu_g^*)$. Then, $(a_{AA}^*, a_{AB}^*, a_{BA}^*, a_{BB}^*, \mu_A^*, \mu_B^*)$ is an EZ-SU.*

Suppose for each g , the extended theory $\bar{\Theta}_g = \mathbb{A}^2 \times \Theta_g$ for some theory Θ_g – that is, each group can make any inference about opponents' strategies. There exists some $\underline{\tau} < 1$ so that for every $\tau \in (\underline{\tau}, 1)$ and $(a_{AA}^, a_{AB}^*, a_{BA}^*, a_{BB}^*, \mu_A^*, \mu_B^*)$ satisfying the above conditions, we have that μ_A^* puts probability 1 on (a_{AA}^*, a_{AB}^*) , μ_B^* puts probability 1 on (a_{BA}^*, a_{BB}^*) , and*

$(a_{AA}^*, a_{AB}^*, a_{BA}^*, a_{BB}^*, \mu_A^*|_{\Theta_A}, \mu_B^*|_{\Theta_B})$ is an EZ, where $\mu_g^*|_{\Theta_g}$ is the marginal of the belief μ_g^* on the theory Θ_g .

Proof. We first consider the case of $\tau = 0$, so the uninformative ex-post signals may be ignored.

For μ a belief and $g \in \{A, B\}$, let $u^\mu(a_i; g)$ represent subjective expected payoff from playing a_i against group g . Suppose $a_{AA}^* \notin \operatorname{argmax}_{\hat{a} \in \mathbb{A}} u^{\mu^*}(\hat{a}; A)$ (the other cases are analogous). By the continuity assumptions on U_A (which is also bounded because Θ_A is bounded), there are some $\epsilon_1, \epsilon_2 > 0$ so that whenever $\mu_i \in \Delta(\bar{\Theta}_A)$ with $\|\mu_i - \mu_A^*\| < \epsilon_1$, we also have $u^{\mu_i}(a_{AA}^*; A) < \max_{\hat{a} \in \mathbb{A}} u^{\mu_i}(\hat{a}; A) - \epsilon_2$. By the definition of asymptotically empirical best responses, find K so that $\mathbf{a}_{A,A}(h)$ must be a myopic ϵ_2 -best response when there are at least K periods of matches against A and B. Agent i has a strictly positive chance to match with groups A and B in every period. So, at all except a null set of points in the probability space, i 's history eventually records at least K periods of play by groups A and B. Also, by assumption, almost surely $\tilde{\mu}_i^{(t)} \rightarrow \mu_A^*$. This shows that by asymptotically myopic best responses, almost surely $\tilde{a}_{iA}^{(k)} \not\rightarrow a_{AA}^*$, a contradiction.

Now suppose some $\theta_A^* = (a_A^*, a_B^*, f^*)$ in the support of μ_A^* does not minimize the weighted KL divergence in the definition of EZ-SU (the case of a model θ_B^* in the support of μ_B^* not minimizing is similar). Then we have

$$\theta_A^* \notin \operatorname{argmin}_{\hat{\theta} \in \bar{\Theta}_A} \left[\begin{array}{l} (\lambda + (1 - \lambda)p_A) \cdot D_{KL}(F^\bullet(a_{AA}^*, a_{AA}^*) \parallel \hat{F}(a_{AA}^*, \hat{a}_A)) \\ + (1 - \lambda)(1 - p_A) \cdot D_{KL}(F^\bullet(a_{AB}^*, a_{BA}^*) \parallel \hat{F}(a_{AB}^*, \hat{a}_B)) \end{array} \right]$$

where $\hat{\theta} = (\hat{a}_A, \hat{a}_B, \hat{F})$.

This is equivalent to:

$$\theta_A^* \notin \operatorname{argmax}_{\hat{\theta} \in \bar{\Theta}_A} \left[\begin{array}{l} (\lambda + (1 - \lambda)p_A) \cdot \mathbb{E}_{y \sim F^\bullet(a_{AA}^*, a_{AA}^*)} \ln(\hat{f}(a_{AA}^*, \hat{a}_A)(y)) \\ + (1 - \lambda)(1 - p_A) \cdot \mathbb{E}_{y \sim F^\bullet(a_{AB}^*, a_{BA}^*)} \ln(\hat{f}(a_{AB}^*, \hat{a}_B)(y)) \end{array} \right]$$

Let this objective, as a function of $\hat{\theta}$, be denoted $WL(\hat{\theta})$. There exists $\theta_A^{opt} = (a_A^{opt}, a_B^{opt}, f^{opt}) \in \bar{\Theta}_A$ and $\delta, \epsilon > 0$ so that $(1 - \delta)WL(\theta_A^{opt}) - 2\delta M - 3\epsilon > (1 - \delta)WL(\theta_A^*)$. By assumption on the primitives, find open neighborhoods V^{opt} and V^* of $\theta_A^{opt}, \theta_A^*$ respectively, so that for all $a_i \in \mathbb{A}$, $g \in \{A, B\}$, $y \in \mathbb{Y}$, $1 - \epsilon \leq f^{opt}(a_i, a_g^{opt})(y) / \hat{f}(a_i, \hat{a}_g)(y) \leq 1 + \epsilon$, for all $\hat{\theta} = (\hat{a}_A, \hat{a}_B, \hat{f}) \in V^{opt}$, and also $1 - \epsilon \leq f^*(a_i, a_g^*)(y) / \hat{f}(a_i, \hat{a}_g)(y) \leq 1 + \epsilon$ for all $\hat{\theta} = (\hat{a}_A, \hat{a}_B, \hat{f}) \in V^*$. Also, by convergence of play in the populations, find T_1 so that in all periods $t \geq T_1$, $\alpha_{AA}^{(t)}(a_{AA}^*) \geq 1 - \delta$ and $\alpha_{BA}^{(t)}(a_{BA}^*) \geq 1 - \delta$.

For $T_2 \geq T_1$, consider a probability space defined by $\Omega := (\{A, B\} \times \mathbb{A}^2 \times (\mathbb{Y})^{\mathbb{A}^2})^\infty$ that describes the randomness in an agent's learning process starting with period $T_2 + 1$. For a

point $\omega \in \Omega$ and each period $T_2 + s$, $s \geq 1$, $\omega_s = (g, a_{-i,A}, a_{-i,B}, (y_{a_i, a_{-i}})_{(a_i, a_{-i}) \in \mathbb{A}^2})$ specifies the group g of the matched opponent, the play $a_{-i,A}, a_{-i,B}$ of hypothetical opponents from groups A and B, and the hypothetical consequence $y_{a_i, a_{-i}}$ that would be generated for every pair of strategies (a_i, a_{-i}) played. As notation, let $opp(\omega, s)$, $a_{-i,A}(\omega, s)$, $a_{-i,B}(\omega, s)$, and $y_{a_i, a_{-i}}(\omega, s)$ denote the corresponding components of ω_s . Define \mathbb{P}_{T_2} over this space in the natural way. That is, it is independent across periods, and within each period, the density (or probability mass function if \mathbb{Y} is finite) of $\omega_s = (g, a_{-i,A}, a_{-i,B}, (y_{a_i, a_{-i}})_{(a_i, a_{-i}) \in \mathbb{A}^2})$ is

$$m_g \cdot \alpha_{AA}^{(T_2+s)}(a_{-i,A}) \alpha_{BA}^{(T_2+s)}(a_{-i,B}) \cdot \prod_{(a_i, a_{-i}) \in \mathbb{A}^2} f^\bullet(a_i, a_{-i})(y_{a_i, a_{-i}}),$$

where m_g is the probability of i from group A being matched up against an opponent of group g , that is $m_A = (\lambda + (1 - \lambda)p_A)$, $m_B = (1 - \lambda)(1 - p_A)$.

For $\theta = (a_A^\theta, a_B^\theta, F^\theta) \in \bar{\Theta}_A$ with f^θ the density of F^θ , $\omega \in \Omega$, consider the stochastic process

$$\ell_s(\theta, \omega) := \frac{1}{s} \sum_{t=T_2+1}^{T_2+s} \ln(f^\theta(a_{AA}^*, a_{opp(\omega, t)}^\theta)(y_{a_{AA}^*, a_{-i, opp(\omega, t)}(\omega, t)}(\omega, t))).$$

By choice of the neighborhood V^* ,

$$\begin{aligned} \limsup_s \sup_{\theta_A \in V^*} \ell_s(\theta_A, \omega) &\leq \epsilon + \frac{1}{s} \sum_{t=T_2+1}^{T_2+s} \ln(f^*(a_{AA}^*, a_{opp(\omega, t)}^*)(y_{a_{AA}^*, a_{-i, opp(\omega, t)}(\omega, t)}(\omega, t))) \\ &\leq \epsilon + \frac{1}{s} \sum_{t=T_2+1}^{T_2+s} \mathbb{1}_{\{a_{-i, opp(\omega, t)}(\omega, t) = a_{opp(\omega, t), A}^*\}} \cdot \ln(f^*(a_{AA}^*, a_{opp(\omega, t)}^*)(y_{a_{AA}^*, a_{opp(\omega, t), A}^*}(\omega, t))) \\ &\quad + (1 - \mathbb{1}_{\{a_{-i, opp(\omega, t)}(\omega, t) = a_{opp(\omega, t), A}^*\}}) \cdot M. \end{aligned}$$

Since $T_2 \geq T_1$, in every period t , $\mathbb{P}_{T_2}(a_{-i, opp(\omega, t)}(\omega, t) = a_{opp(\omega, t), A}^*) \geq 1 - \delta$. Let $(\xi_k)_{k \geq 1}$ a related stochastic process: it is i.i.d. such that each ξ_k has δ chance to be equal to M , $(1 - \delta)m_A$ chance to be distributed according to $\ln(f^*(a_{AA}^*, a_A^*)(y))$ where $y \sim f^\bullet(a_{AA}^*, a_{AA}^*)$, and $(1 - \delta)m_B$ chance to be distributed according to $\ln(f^*(a_{AB}^*, a_B^*)(y))$ where $y \sim f^\bullet(a_{AB}^*, a_{BA}^*)$. By law of large numbers, $\frac{1}{s} \sum_{k=1}^s \xi_k$ converges almost surely to $\delta M + (1 - \delta)WL(\theta_A^*)$. By this comparison, $\limsup_s \sup_{\theta_A \in V^*} \ell_s(\theta_A, \omega) \leq \epsilon + \delta M + (1 - \delta)WL(\theta_A^*)$ \mathbb{P}_{T_2} -almost surely. By a similar argument, $\liminf_s \inf_{\theta_A \in V^{opt}} \ell_s(\theta_A, \omega) \geq -\epsilon - \delta M + (1 - \delta)WL(\theta_A^{opt})$ \mathbb{P}_{T_2} -almost surely.

Along any ω where we have both $\limsup_s \sup_{\theta_A \in V^*} \ell_s(\theta_A, \omega) \leq \epsilon + \delta M + (1 - \delta)WL(\theta_A^*)$ and $\liminf_s \inf_{\theta_A \in V^{opt}} \ell_s(\theta_A, \omega) \geq -\epsilon - \delta M + (1 - \delta)WL(\theta_A^{opt})$, if ω also leads to i always playing a_{AA}^* against group A and a_{AB}^* against group B in all periods starting with $T_2 + 1$, then the posterior belief assigns to V^* must tend to 0, hence $\tilde{\mu}_i^{(t)} \not\rightarrow \mu_A^*$. Starting from any length T_2 history h , there exists a subset $\hat{\Omega}_h \subseteq \Omega$ that leads to i not playing the EZ-SU strategy in at least one period starting with $T_2 + 1$. So conditional on h , the probability

of $\tilde{\mu}_i^{(t)} \rightarrow \mu_A^*$ is no larger than $1 - \mathbb{P}_{T_2}(\hat{\Omega}_h)$. The unconditional probability is therefore no larger than $\mathbb{E}_h[1 - \mathbb{P}_{T_2}(\hat{\Omega}_h)]$, where \mathbb{E}_h is taken with respect to the distribution of period T_2 histories for i . But this term is also the probability of i playing non-EZ-SU action at least once starting with period T_2 . Since there are finitely many actions and $(\tilde{a}_{iA}^{(t)}, \tilde{a}_{iB}^{(t)}) \rightarrow (a_{AA}^*, a_{AB}^*)$ almost surely, $\mathbb{E}_h[1 - \mathbb{P}_{T_2}(\hat{\Omega}_h)]$ tends to 0 as $T_2 \rightarrow \infty$. We have a contradiction as this shows $\tilde{\mu}_i^{(t)} \not\rightarrow \mu_A^*$ with probability 1.

Now consider the foundation for EZs. Suppose Let $\bar{K} < \infty$ be an upper bound on $K_{g,g'}(a_i, a_{-i}; (a_A, a_B, F))$ across all $g, g' \in \{A, B\}$, $a_i, a_{-i} \in \mathbb{A}$, $(a_A, a_B, F) \in \bar{\Theta}_g$. Here \bar{K} is finite because \mathbb{A} is finite and $K_{g,g'}$ is continuous in the extended model, which is from a compact domain. Let $F_\tau^X(a_{-i}) \in \Delta(\mathbb{A})$ represent the distribution of ex-post signals given precision τ , when opponent plays $a_{-i} \in \mathbb{A}$. It is clear that there exists some $\underline{\tau} < 1$ so that for any $a_{-i} \neq a'_{-i}$, $\tau \in (\underline{\tau}, 1)$, we get $\min(m_A, m_B) \cdot D_{KL}(F_\tau^X(a_{-i}) \parallel F_\tau^X(a'_{-i})) > \bar{K}$. Therefore, given any $(a_{AA}^*, a_{AB}^*, a_{BA}^*) \in \mathbb{A}^3$, the solution to

$$\min_{\hat{\theta} \in \bar{\Theta}_A} \left[\begin{aligned} & (\lambda + (1 - \lambda)p_A) \cdot [D_{KL}(F^\bullet(a_{AA}^*, a_{AA}^*) \parallel \hat{F}(a_{AA}^*, \hat{a}_A)) + D_{KL}(F_\tau^X(a_{AA}^*) \parallel F_\tau^X(\hat{a}_A))] \\ & + (1 - \lambda)(1 - p_A) \cdot [D_{KL}(F^\bullet(a_{AB}^*, a_{BA}^*) \parallel \hat{F}(a_{AB}^*, \hat{a}_B)) + D_{KL}(F_\tau^X(a_{BA}^*) \parallel F_\tau^X(\hat{a}_B))] \end{aligned} \right]$$

must satisfy $\hat{a}_A = a_{AA}^*$, $\hat{a}_B = a_{BA}^*$, because (a_{AA}^*, a_{BA}^*, F) for any $F \in \Theta_A$ has a KL divergence no larger than \bar{K} . On the other hand, any $(\hat{a}_A, \hat{a}_B, \hat{F})$ with either $\hat{a}_A \neq a_{AA}^*$ or $\hat{a}_B \neq a_{BA}^*$ has KL divergence strictly larger than \bar{K} by the choice of τ . The rest of the argument is similar to the case of EZ-SU. \square