Higher-Order Beliefs and (Mis)learning from Prices*

Kevin He† Jonathan Libgober‡

May 10, 2025

Abstract

We study misperceptions of private-signal correlation when agents repeatedly match up to play an incomplete-information Cournot duopoly game. We find that a misperception's viability can depend on whether agents hold flexible or dogmatically correct beliefs about price elasticity. If agents have flexible beliefs and learn elasticity by observing prices, correlation misperceptions indirectly distort behavior through elasticity misinference. If agents know the true elasticity, this learning channel is eliminated. Correlation misperceptions have opposite direct and indirect effects on behavior, so the presence of elasticity inference can reverse an error's viability.

*Results from this paper originally appeared in "Evolutionarily Stable (Mis)specifications: Theory and Applications." We thank Cuimin Ba, Thomas Chaney, Sylvain Chassang, In-Koo Cho, Krishna Dasaratha, Andrew Ellis, Ignacio Esponda, Mira Frick, Drew Fudenberg, Alice Gindin, Ryota Iijima, Yuhta Ishii, Philippe Jehiel, Pablo Kurlat, Elliot Lipnowski, Jonny Newton, Filippo Massari, Andy Postlewaite, Luis Rayo, Philipp Sadowski, Alvaro Sandroni, Grant Schoenebeck, Joshua Schwartzstein, Philipp Strack, Carl Veller, and various conference and seminar participants for helpful comments. Byunghoon Kim provided excellent research assistance. Kevin He thanks the California Institute of Technology for hospitality when some of the work on this paper was completed, and the University Research Foundation Grant at the University of Pennsylvania for financial support. Jonathan Libgober thanks Yale University and the Cowles Foundation for their hospitality.

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

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

1 Introduction

In strategic situations where players face uncertainty over the state of nature, agents' behavior can depend on both their beliefs about the state (i.e., first-order beliefs) and their beliefs about other players' information (i.e., higher-order beliefs). But, significant evidence suggests economic actors often find it difficult to form accurate higher-order beliefs or detect systematic biases in them. This paper investigates mistaken higher-order beliefs from an evolutionary perspective, asking which errors might confer an advantage and when.

Our main message is that whether a given misperception in higher-order beliefs improves or harms payoffs can depend on whether other, persistent parameters of the game are known or inferred. A higher-order misperception can influence the agent's conjecture of opponent behavior, since the opponent's action conditions on their information. This distorted conjecture can directly distort the agent's action. But more subtly, when the agent does not know the true values of the persistent game parameters and must infer them through repeated play, the same mispredictions about opponents' behavior cause the agent to misinterpret game outcomes. This misinference induces distorted beliefs about the game parameters, possibly letting the agent commit to strategically beneficial behavior. So, in addition to its direct effect, higher-order misperceptions can distort behavior indirectly through this *learning channel*. Given that errors can have opposite direct and indirect effects on behavior, an agent's knowledge or ignorance about the persistent parameters (and thus whether the learning channel is present) can determine whether the error facilitates such beneficial commitments.

1.1 Summary of the Setup and Main Results

We illustrate this idea in the context of a linear-quadratic-normal (LQN) Cournot duopoly game of incomplete information, similar to Vives (1988).[1] The state is the intercept of the demand curve (i.e., demand shock), drawn i.i.d. from a normal distribution each time the game is played, with players receiving possibly correlated signals before choosing quantities. The key persistent game parameter is the slope of the demand curve (i.e., price elasticity), which players may or may not know. Setups within the LQN family have received significant attention in part because they admit tractable comparative statics with respect to players' information (illustrated in Bergemann and Morris (2013), as well as Miyashita and Ui (2023);

[1] Angeletos and Pavan (2007) extend this model to more general environments. We explain how our results generalize in Section 5.3.

Bergemann, Heumann, and Morris (2017)). Here, we use this setup to study misperceptions of others' information.

Society consists of *residents* and *entrants* who repeatedly pair up to play the LQN game. Price elasticity is a structural parameter of the market that remains fixed across games, but idiosyncratic daily demand shocks cause the demand intercept to be drawn i.i.d. across games. Residents correctly know how the demand signals of the two participants in each game are correlated, whereas entrants may misperceive this correlation in information. We show that if agents correctly know the persistent price elasticity, then assortative matching (i.e., entrants' welfare is determined by how they do when playing against each other) favors entrants who overestimate correlation in different players' signals. On the other hand, uniform matching (i.e., entrants' welfare is determined by how they do when playing against residents) favors entrants who underestimate said correlation. However, the situation is exactly reversed when agents do not know price elasticity and infer this parameter from game outcomes. That is, sometimes erroneous higher-order beliefs only benefit agents who are uncertain about the game's persistent parameters.

To see the intuition behind this finding, consider an agent who misperceives signals to be excessively correlated — an error we refer to as *projection bias*. We show that the welfare implications of projection bias with uniform matching depend on whether the bias induces more aggressive strategies in equilibrium — that is, strategies that respond more to changes in private information about demand. Using a more aggressive strategy acts as a commitment that induces the opponent to behave less aggressively, which is beneficial as this game features strategic substitutes. Thus, we examine whether projection bias increases the aggressiveness of subjective best responses.

On the one hand, the direct effect of projection bias makes agents act less aggressively. When an agent has a private signal that suggests high market demand, they overestimate the similarity of their opponent's information and thus exaggerate how much the other player will increase their production level. This force limits how much the agent wishes to increase production, since the two competitors' quantity choices are strategic substitutes. The bias thus harms the agent's profits when they know price elasticity.

On the other hand, the indirect effect of projection bias through the learning channel acts in the opposite direction. Suppose this agent infers elasticity from prices. Then, projection bias causes the agent to underestimate price elasticity. This is because the agent's bias leads them to overestimate the correlation between their own signal realization and their opponent's production quantity in each game. After a high private signal, the market price remains higher than the agent expects, which they rationalize by inferring a low price elasticity. Underinferring price elasticity increases the aggressiveness of the agent's best response, as they underestimate how quickly the price decreases when they produce more. Thus, the learning channel influences the aggressiveness of the strategy in precisely the opposite way compared to the misperception of the signal correlation.

While this observation may suggest that the overall impact of correlation misperception is ambiguous, we show that the indirect effect is in fact stronger than the direct effect. Intuitively, this is because elasticity influences strategies much more than perceived signal correlation. However, the indirect effect is present only when agents are initially uncertain about price elasticity. Putting everything together, we conclude that projection bias can only invade a rational society when the entrants draw inferences about price elasticity from market prices, not when they already know price elasticity with certainty. (Our results for assortative matching are essentially the inverse of these findings, as assortative matching favors *less* aggressive strategies rather than more aggressive ones.)

1.2 Related Literature

The question of whether profit maximization requires firms to behave (at least as-if) rationally has been of interest to economists since at least Friedman (1953)'s market-selection hypothesis. The subsequent literature points out that while misperceptions can only lower payoffs in decision problems, the same need not be true in strategic settings as a firm's performance also depends on how its competitors react to it. To study the consequences of the market selection pressures on errors in higher-order beliefs, our work applies stability concepts from the literature on the *indirect evolutionary approach* (surveyed in Alger and Weibull (2019); Robson and Samuelson (2011)) to these biases.

Typically, the indirect evolutionary approach assumes agents in a society are endowed with different subjective preferences over game outcomes. If a "new preference" leads to higher objective payoffs in equilibrium than an "existing preference" when the latter is dominant in the society, then we say the former has an evolutionary advantage and can invade the latter. In our application, a higher-order misperception is equivalent to a subjective preference only if the agent knows the persistent game parameters. If the agent is uncertain about these parameters and infers them from game outcomes, then the same misperception can lead to different beliefs about these parameters when the environment varies, as highlighted in our companion paper He and Libgober (2024). A recent literature on misspecified Bayesian learning (Esponda and Pouzo (2016); Frick, Iijima, and Ishii (2024); Heidhues, Koszegi, and Strack (2018), among others) studies the implications of mislearning persistent parameters of the environment on behavior and welfare. Our contribution is to study the equilibrium consequences of misperceiving others' information within a seminal game, and to ask how the answer depends on agents' uncertainty about the persistent game parameters. Our interest in matching assortativity is borrowed from Alger and Weibull (2013), who study the effect of assortativity on the selection of preferences and find that more assortativity selects for less selfish behavior. Alger and Weibull (2019) discuss interest in this aspect of the matching process in depth, as assortativity is naturally generated by various population structures.

One theme in this literature is that rational payoff maximization cannot be evolutionarily suboptimal unless agents' preferences or strategies are at least partially observable by others. Dekel, Ely, and Yilankaya (2007) characterize stable preferences in two-by-two games under the assumption of observable preferences, while also showing that rational preferences are favored when preferences are unobservable. Heifetz, Shannon, and Spiegel (2007) show that distortions are evolutionarily beneficial within a general framework that allows richer action spaces, again under the assumption of observable preferences (although their conclusions may remain even when this assumption is relaxed). Our framework analogously assumes that agents' perceptions of signal correlation are observable, which provides scope for departures from rationality (as in these other works). Our results also rely on the observability of strategies; we discuss the implications of relaxing this assumption in Section 5.1.

Much of the past work using the evolutionary approach focuses on stage games with complete information. By contrast, we study a stage game with incomplete information since we are interested in biases that involve misperceptions of others' information. This presents additional challenges for characterizing equilibrium, since the game's strategy space becomes much richer when players can condition their actions on their private signals. Other papers in economic theory have studied the implications of information projection or the related bias of taste projection (Gagnon-Bartsch, Pagnozzi, and Rosato (2021); Gagnon-Bartsch and Rosato (2024); Madarász (2012)). But, we pinpoint a novel mechanism where a misperception of correlation in information grants a strategic advantage by causing the agent to mislearn some persistent parameter of the game through repeated play, and thus commit to a more beneficial strategy.

The idea that firms may be misspecified relates to a line of work studying pricing algorithms. These papers usually consider an environment where each firm uses a learning algorithm to estimate its profit function. A common theme is that if the learning algorithm is misspecified (as it often must be given the complexity of the market environment), then the prices algorithms converge to can end up being excessively collusive. For example, Hansen, Misra, and Pai (2021); Asker, Fershtman, and Pakes (2023); Calvano, Calzolari, Denicoló, and Pastorello (2020) use numerical simulations to study reinforcement-learning algorithms that assume the firm is facing a time-stationary competitive environment, when in reality they face competition from other learning algorithms that adjust their behavior over time. In our setting, we can interpret an agent's misperceptions of signal correlation as a misspecification encoded in the agent's pricing algorithm, and our results similarly show that Bayesian algorithms that estimate market price elasticity under misspecifications can end up behaving too cooperatively or too aggressively (depending on the error). Also related to our work is Berman and Heller (2024), who consider firms that choose from a broad class of possibly non-Bayesian learning algorithms. By comparison, we are closer to the misspecified Bayesian learning literature as we restrict attention to only agents/algorithms that draw Bayesian inferences given their misperceptions.

2 Framework

Following the indirect evolutionary approach and our companion paper He and Libgober (2024), we study an environment where a continuum of agents are matched up in pairs each period to play a two-player stage game.

2.1 Stage Game and Information Structure

We begin by describing the stage game, a simultaneous-move game with incomplete information. There is a demand state $\omega \sim \mathcal{N}(0, \sigma_\omega^2)$, where $\mathcal{N}(\mu, \sigma^2)$ is the normal distribution with mean μ and variance σ^2. Firm i observes a private signal $s_i = \omega + \epsilon_i$, and then chooses a quantity $q_i \in \mathbb{R}$. The resulting market price is $P = \omega - r^\bullet \cdot \frac{1}{2}(q_1 + q_2) + \zeta$, where $\zeta \sim \mathcal{N}(0, (\sigma_\zeta^\bullet)^2)$ is a price shock independent of other random variables. Throughout the paper, we use superscript $^\bullet$ to denote the true parameters, distinguishing them from the subjectively believed parameters which we describe below. The firm pays a cost $\frac{1}{2}q_i^2$ when it chooses quantity q_i, which leads to a profit of $q_i P - \frac{1}{2} q_i^2$ if the market price is P.

As in many other LQN oligopoly models, 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 the market price decreases in total supply. When $P < 0$, the market pays for disposal. 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 allow players' signals in a given stage game to be correlated conditional on ω. We study misperceptions of this correlation. Recalling that $s_i = \omega + \epsilon_i$, we take:

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

where $\eta_i \sim \mathcal{N}(0, \sigma_\epsilon^2)$ is the idiosyncratic component generated i.i.d. across players and $z \sim \mathcal{N}(0, \sigma_\epsilon^2)$ is the common component. Under this parameterization, $\kappa^\bullet$ reflects the similarity between the players' private information. In particular, higher $\kappa^\bullet$ leads to an information structure with higher conditional correlation. Indeed, when $\kappa^\bullet = 0$, s_i and s_{-i} are conditionally uncorrelated given ω. On the other hand, when $\kappa^\bullet = 1$, we always have $s_i = s_{-i}$ (i.e., perfect correlation between signals). Our functional form for ϵ_i ensures $\mathrm{Var}(s_i)$ is constant in $\kappa^\bullet$, so that the distribution of ω given s_i does not vary with $\kappa^\bullet$.

As mentioned, we embed this stage game within a larger framework to discuss the selection of misspecified models. This elaboration takes the persistent parameters of the stage game to be $\sigma_\omega^2 > 0$ (variance of demand state), $r^\bullet > 0$ (a measure of the elasticity of market price with respect to quantity supplied), $(\sigma_\zeta^\bullet)^2 > 0$ (variance of price shock), and $\kappa^\bullet \in [0,1]$ (a measure of signal correlation). By contrast, each time the stage game is played, ω, z, η_i and ζ are independently drawn from their respective distributions.

2.2 Models, Inference, and Strategies

The stage game is common knowledge except for the persistent parameters $\kappa^\bullet, r^\bullet$, and $(\sigma_\zeta^\bullet)^2$. Agents interpret their environment through their *models* of the world. A model can have two kinds of parameters: *free parameters* are estimated using game outcomes, while *fixed parameters* are dogmatically given by the model and not subject to inference. Signal correlation is a fixed parameter in every model, so different models can encode different dogmatic beliefs about that aspect of the stage game. We consider both *flexible* models where signal correlation $\tilde{\kappa}$ is a fixed parameter but price elasticity $\tilde{r}$ and price shock variance $\tilde{\sigma}_\zeta^2$ are free parameters,[2] as well as *dogmatic* models where $\tilde{\kappa}, \tilde{r}, \tilde{\sigma}_\zeta^2$ are all fixed parameters.

Our interest will be in studying misperceptions of signal correlation. We distinguish between the two possible directions of this misperception:

Definition 1. Let $\tilde{\kappa}$ be a player's perceived κ. A player exhibits *correlation neglect* if $\tilde{\kappa} < \kappa^\bullet$. A player exhibits *projection bias* if $\tilde{\kappa} > \kappa^\bullet$.

Correlation neglect agents underestimate the correlation between players' signals in the stage game, whereas projection bias agents exaggerate this correlation. We are agnostic about the origin of these misspecifications, except to say that in many contexts they do seem to arise, as highlighted in our discussion in Section 1.2. However, our interest in this paper is whether misspecifications of this form can invade a rational society.

We now describe inference for flexible models. A *consequence* is a triple (s_i, q_i, P) that contains i's signal, i's quantity choice, and the realized market price. A *strategy* for i is a quantity choice as a function of i's signal realization, $Q_i(s_i)$. Let $\mathbb{Y}$ denote the set of all consequences, and let $\mathbb{S}$ denote the space of strategies. For each $(\kappa, r, \sigma_\zeta^2)$, we define $F_{\kappa,r,\sigma_\zeta^2} : \mathbb{S} \times \mathbb{S} \to \Delta(\mathbb{Y})$ to be the mapping between strategy profiles and the distribution over i's consequences in a stage game with parameters $(\kappa, r, \sigma_\zeta^2)$. The following definition captures our notion of free-parameter estimation. This inference is performed as a function of a given stage-game strategy profile:

Definition 2. Let $F^\bullet(Q_i, Q_{-i})$ denote the objective distribution over i's consequences given strategy profile Q_i, Q_{-i}. We say that inference $(\tilde{r}, \tilde{\sigma}_\zeta^2)$ is a *self-confirming inference given strategy profile* Q_i, Q_{-i} *and correlation* κ if $F^\bullet(Q_i, Q_{-i}) = F_{\kappa,\tilde{r},\tilde{\sigma}_\zeta^2}(Q_i, Q_{-i})$.

Self-confirming inferences are not falsified by the distribution of consequences that a player sees when they perceive correlation κ and repeatedly play the stage game using strategy Q_i against different opponents who all use the strategy Q_{-i}. Self-confirming inferences need not exist in general, in which case a goodness-of-fit criterion would be necessary for inferences to be well-defined. Esponda and Pouzo (2016) motivate KL-divergence as a natural criterion for

[2]While a flexible model allows agents to infer both r and σ_ζ^2, their misinference about r drives the results. Since each player's profit is linear in the market price, belief about the variance of the idiosyncratic price shock does not change their expected payoffs or behavior. The parameter σ_ζ^2 absorbs changes in the variance of market price, creating significant tractability.

misspecified Bayesian agents. However, to avoid complications, our analysis below will focus on values of the true parameters such that self-confirming inferences exist.

Next, we present a partial equilibrium notion where both players choose strategies that maximize profit given their beliefs about the persistent parameters, and said beliefs for some player i are either the fixed parameters $\tilde{\kappa}, \tilde{r}, \tilde{\sigma}_\zeta^2$ (if i has a dogmatic model) or fixed parameter $\tilde{\kappa}$ together with the self-confirming inferences (if i has a flexible model). The reason why we do not require $-i$ to also derive beliefs from the same interaction is that some of our analysis concerns environments where $-i$'s beliefs are primarily shaped by the consequences they observe in other matches—in particular, when i is part of a negligible subpopulation.

Definition 3. A strategy profile Q_i, Q_{-i} and belief profile $(\tilde{\kappa}_i, \tilde{r}_i, \tilde{\sigma}^2_{\zeta,i})$, $(\tilde{\kappa}_{-i}, \tilde{r}_{-i}, \tilde{\sigma}^2_{\zeta,-i})$ are a *linear partial equilibrium* if

- For each player k, $Q_k(s_k) = \alpha_k s_k$ for some $\alpha_k \geq 0$.
- For each player k, Q_k is an interim-stage best response against the opponent's strategy given belief $(\tilde{\kappa}_k, \tilde{r}_k, \tilde{\sigma}^2_{\zeta,k})$.
- For the first player i, $\tilde{\kappa}_i$ is the fixed parameter given by i's model, and $(\tilde{r}_i, \tilde{\sigma}^2_{\zeta,i})$ are either the fixed parameters given by i's dogmatic model or i's self-confirming inference given Q_i, Q_{-i}, and $\tilde{\kappa}_i$ (when i has a flexible model).

This definition reflects *partial* equilibrium since we only restrict the inferences of the first player and not the second player. Our focus on linear strategies follows other work studying LQN games. Since the best response (among the family of all strategies) to any linear strategy is linear for any belief about the correlation parameter and price elasticity (shown in Lemma 2), we focus on equilibria where everyone uses linear strategies. We sometimes refer to the linear strategy $s_i \mapsto \alpha_i s_i$ simply as α_i.

2.3 Stability and Invasion

Our analysis will compare the *entrant model*, which is used by an infinitesimally small group of *entrants* in the population, with the *resident model*, which is used by the remaining group called the *residents*. We will consider two particular interaction structures.[3] In uniform

[3]Appendix B considers a more general formulation which includes these two interaction structures as extreme cases.

matching, each agent is matched with an opponent drawn uniformly at random from the entire population, with agents observing their opponent's model. So, agents are only matched against the entrant group with infinitesimal probability. In assortative matching, agents are always matched within the group that uses the same model.[4]

In what follows, we use the subscript R to refer to the resident and the subscript E to refer to the entrant. For example, κ_R denotes the resident's perceived correlation parameter, and κ_E denotes that of the entrant. We let $\alpha_{g\to g'}$ denote the strategy that a group g agent uses when matched against someone from group g'. For strategies α_g, α_{-g} in the stage game, let $U^\bullet(\alpha_g, \alpha_{-g})$ be the objective expected utility of playing strategy α_g against α_{-g}. We refer to the objective expected utility of agents who use a model as that model's *fitness*.

Definition 4. With uniform matching, a *linear equilibrium* consists of strategies $\alpha_{R\to R}, \alpha_{R\to E}, \alpha_{E\to R}$ and beliefs $(\tilde{\kappa}_R, \tilde{r}_R, \tilde{\sigma}^2_{\zeta,R})$, $(\tilde{\kappa}_E, \tilde{r}_E, \tilde{\sigma}^2_{\zeta,E})$ such that:

- $\alpha_{R\to R}, \alpha_{R\to R}, (\tilde{\kappa}_R, \tilde{r}_R, \tilde{\sigma}^2_{\zeta,R}), (\tilde{\kappa}_R, \tilde{r}_R, \tilde{\sigma}^2_{\zeta,R})$ are a linear partial equilibrium,
- $\alpha_{E\to R}, \alpha_{R\to E}, (\tilde{\kappa}_E, \tilde{r}_E, \tilde{\sigma}^2_{\zeta,E}), (\tilde{\kappa}_R, \tilde{r}_R, \tilde{\sigma}^2_{\zeta,R})$ are a linear partial equilibrium.

We say κ_R is *resistant to invasion from* κ_E *with uniform matching* if $U^\bullet(\alpha_{R\to R}, \alpha_{R\to R}) \geq U^\bullet(\alpha_{E\to R}, \alpha_{R\to E})$ in every linear equilibrium, and we say κ_R is *susceptible to invasion with uniform matching* if $U^\bullet(\alpha_{R\to R}, \alpha_{R\to R}) < U^\bullet(\alpha_{E\to R}, \alpha_{R\to E})$ in every linear equilibrium.

With assortative matching, a *linear equilibrium* consists of strategies $\alpha_{R\to R}, \alpha_{E\to E}$ and beliefs $(\tilde{\kappa}_R, \tilde{r}_R, \tilde{\sigma}^2_{\zeta,R})$, $(\tilde{\kappa}_E, \tilde{r}_E, \tilde{\sigma}^2_{\zeta,E})$ such that:

- $\alpha_{R\to R}, \alpha_{R\to R}, (\tilde{\kappa}_R, \tilde{r}_R, \tilde{\sigma}^2_{\zeta,R}), (\tilde{\kappa}_R, \tilde{r}_R, \tilde{\sigma}^2_{\zeta,R})$ are a linear partial equilibrium,
- $\alpha_{E\to E}, \alpha_{E\to E}, (\tilde{\kappa}_E, \tilde{r}_E, \tilde{\sigma}^2_{\zeta,E}), (\tilde{\kappa}_E, \tilde{r}_E, \tilde{\sigma}^2_{\zeta,E})$ are a linear partial equilibrium.

We say κ_R is *resistant to invasion from* κ_E *with assortative matching* if $U^\bullet(\alpha_{R\to R}, \alpha_{R\to R}) \geq U^\bullet(\alpha_{E\to E}, \alpha_{E\to E})$ in every linear equilibrium, and we say κ_R is *susceptible to invasion with assortative matching* if $U^\bullet(\alpha_{R\to R}, \alpha_{R\to R}) < U^\bullet(\alpha_{E\to E}, \alpha_{E\to E})$ in every linear equilibrium.

This definition embeds the idea that agents with flexible models correctly think that the values of the persistent game parameters do not change depending on the group membership

[4]Our assumption that entrants form an infinitesimally small group, as in He and Libgober (2024), is made to avoid technical complications. Appendix B explains in detail how to extend our baseline framework to accommodate the possibility that the entrants forms a very small but positive-mass group. We note that this modification does not meaningfully change any of our results.

of the opponent. In particular, in an environment with uniform matching and with residents who have flexible models, the residents' beliefs about the free parameters $(\tilde{r}_R, \tilde{\sigma}^2_{\zeta,R})$ when playing against entrants are estimated using the consequences in their matches against other residents. This feature arises because these residents use all their available data to estimate the persistent game parameters, and matches against entrants comprise an infinitesimally small portion of their data.

2.4 Discussion of the Framework

In our framework, each model is a parametric class of data-generating processes. Misspecified models are often used in the misspecified learning literature to represent and study behavioral biases – in our case, errors in beliefs about signal correlation in the stage game. Our subsequent results will compare the payoff implications of different models with different perceptions of signal correlation, and sometimes focus on comparing entrant models that are "close" to the resident models in that they only differ slightly in this perception. This is in line with some of the recent work on misspecified learning (e.g., Fudenberg and Lanzani (2023)) but complementary to other work on preference evolution that instead focuses on analyzing stability against the universe of all possible subjective preferences, typically in 2-by-2 games.

We find it worthwhile to study errors in higher-order beliefs both because these biases have received significant attention recently in the behavioral literature (e.g., Gagnon-Bartsch et al. (2021); Gagnon-Bartsch and Rosato (2024); Madarász (2012)) and because they are likely harder to detect than errors in first-order beliefs and thus more likely to emerge in the first place. We formalize one version of the claim that higher-order errors are "harder to detect" in Section 5.2.2. We show that in our setting, the higher-order errors we study cause no loss in profit and no misinference of the persistent parameters when agents act as monopolists. By contrast, first-order errors in beliefs about the slope or intercept of the demand curve lead to losses when agents are monopolists.

As is typical in the indirect evolutionary approach literature, our solution concept of linear equilibrium assumes players know the strategies used by different groups of agents in equilibrium.[5] In our setting, we could justify such an assumption by imagining that agents sample a large number of others from each of the resident and entrant populations and observe

[5]Even papers that consider imperfect observability of other's types assume that agents correctly know the equilibrium mapping from types to strategies (Dekel, Ely, and Yilankaya, 2007).

their signals and quantities, thus learning how signals map into quantity choices in each population. Our results on the selection of biases depend critically on this assumption. Section 5.1 shows that if the misspecified entrants can additionally hold a dogmatic misperception of others' strategies, then *every* misperception of signal correlation is equally viable: they can all lead to the Stackelberg payoff[6] in the stage game.

Finally, implicit in our definition of resistance and susceptibility to invasion is the idea that agents interact frequently enough as to settle into a linear equilibrium, and that the payoffs in the linear equilibrium determine the viability of biases. This is in line with the existing work on the indirect evolutionary approach, which usually uses equilibrium payoffs to evaluate the long-run fitness of different types (see, for instance, Dekel et al. (2007); Alger and Weibull (2019)). A typical justification for equilibrium play is based on the idea that play settles on equilibrium more quickly than the timescale of evolution (e.g., footnote 10 of Dekel et al. (2007)).

3 Subjective Best Responses and Self-Confirming Inferences

This section presents results characterizing best responses and self-confirming inferences. These preliminary lemmas enable an explicit description of equilibrium outcomes which we will subsequently apply to discuss the evolutionary selection of models. Our first result shows that when i sees private signal s_i, their mean posterior beliefs about the state and about opponent's signal are linear functions of s_i.

Lemma 1. *There exists a strictly increasing function $\psi(\kappa)$, with $\psi(0) > 0$ and $\psi(1) = 1$, so that:*

$$\mathbb{E}_\kappa[s_{-i} \mid s_i] = \psi(\kappa) \cdot s_i, \text{ for all } s_i \in \mathbb{R} \text{ and } \kappa \in [0, 1].$$

In addition, there exists a strictly positive constant $\gamma > 0$ so that

$$\mathbb{E}_\kappa[\omega \mid s_i] = \gamma \cdot s_i, \text{ for all } s_i \in \mathbb{R}, \kappa \in [0, 1].$$

This result uses the tractability of the LQN framework. The coefficient γ that characterizes

[6] The Stackelberg payoffs corresponds to the player choosing a linear strategy maximizing their payoff, subject to the constraint that the opponent plays a best reply to the strategy.

an agent's inference about the state does not depend on their perception of κ. But higher κ implies the agent infers more about the opponent's signal from their signal. In other words, a misperception of κ only distorts the agent's higher-order belief about the opponent's signal realization (and hence, the opponent's belief), but does not affect the agent's first-order belief about the state ω.

The linearity of $\mathbb{E}[\omega \mid s_i]$ and $\mathbb{E}[s_{-i} \mid s_i]$ in s_i provided by Lemma 1 gives us an explicit characterization of best responses in the stage game, given beliefs about the κ and r parameters. Specifically, Lemma 1 implies that the expected price given s_i is a linear function of s_i when the opponent follows a linear strategy. Our next lemma uses this fact to express player i's expected payoff as a quadratic function of α_i. In what follows, we let $U_i(\alpha_i, \alpha_{-i}; \kappa, r)$ denote the subjective expected profit of player i who perceives correlation parameter κ and believes elasticity to be r, when playing strategy α_i and facing strategy α_{-i}:

Lemma 2. *For linear strategies α_i, α_{-i} and perceived parameters $\kappa \in [0,1], r \geq 0$, we have:*

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

For the same parameters, the linear strategy

$$\alpha_i^{BR}(\alpha_{-i}; \kappa, r) := \frac{\gamma - \frac{1}{2} r \psi(\kappa) \alpha_{-i}}{1 + r}$$

subjectively best responds to α_{-i} at the interim stage among all (possibly non-linear) strategies $Q_i : \mathbb{R} \to \mathbb{R}$.

One key insight of Lemma 2 is that an agent's subjective expected utility and subjective best response depend on their beliefs about κ and r, but not σ_ζ^2. Call a linear strategy more *aggressive* if its coefficient $\alpha_i \geq 0$ is larger. Lemma 2 implies that agent i's subjective best response function becomes more aggressive when i believes in lower κ or lower r. The intuition for this was outlined in the introduction. We have $\frac{\partial \alpha_i^{BR}}{\partial \kappa} < 0$ as the agent can better leverage 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 inelastic demand induces the agent to behave more aggressively, since prices become less responsive to quantity choices.

Lemma 2 calculates the subjective expected utility, and uses this expression to determine the best response given these perceptions. However, an immediate corollary is that the *objective* welfare coincides with this expression evaluated at $r = r^\bullet$ and $\kappa = \kappa^\bullet$; that is,

$$\mathbb{E}[s_i^2] \cdot \left(\alpha_i \gamma - \frac{1}{2} r^\bullet \alpha_i^2 - \frac{1}{2} r^\bullet \psi(\kappa^\bullet) \alpha_i \alpha_{-i} - \frac{1}{2} \alpha_i^2\right)$$

This observation is useful for our fitness calculations below, where objective welfare and perceived welfare may differ.

Finally, we characterize self-confirming inference given a strategy profile and a correlation perception.

Lemma 3. *There exists some $L > 0$ such that a unique self-confirming inference exists for any $\kappa \in [0, 1]$ and $0 \leq \alpha_i, \alpha_{-i} \leq \gamma$ whenever $(\sigma_\zeta^\bullet)^2 \geq L$. In this case, the self-confirming inference for elasticity is*

$$r_i^{INF}(\alpha_i, \alpha_{-i}, ; \kappa^\bullet, \kappa, r^\bullet) := r^\bullet \frac{\alpha_i + \alpha_{-i}\psi(\kappa^\bullet)}{\alpha_i + \alpha_{-i}\psi(\kappa)}.$$

Lemma 3 shows that for agents with flexible models, there is a unique inference of the free parameters r, σ_ζ^2 that perfectly matches the observed price distribution for any linear strategy profile, provided the true price shock variance is large enough and both agents' strategies are less aggressive than γ. Note that by Lemma 2, i's best response against any α_{-i} is always bounded by γ, given any beliefs $\kappa \in [0, 1], r \geq 0$. Therefore, no linear equilibrium exists where either player uses a strategy $\alpha_i > \gamma$ and hence we do not need to worry about this restriction when we compute equilibrium strategies.

The self-confirming property always holds in the linear equilibria we use to define resistance to invasion. In the proof of Lemma 3, we define L as the largest price variance possible when $\zeta = 0$, for any κ and any strategies less aggressive than γ. As players change strategies, the corresponding variance of the price will change as well. By imposing a lower bound on $(\sigma_\zeta^\bullet)^2$, no matter what inference or strategy emerges in the equilibrium, players can infer σ_ζ^2 to match the variance perfectly, and in particular they can do so independently of their inference about the mean of the price distribution. This observation implies that the self-confirming inference about r is the unique one such that the expected mean of the price distribution matches the actual price distribution. While the value of L we define in the proof is larger than necessary to ensure that a self-confirming inference exists, it allows us to avoid placing joint restrictions on parameters and equilibrium strategies.

A key lesson of Lemma 3 is that for a fixed strategy profile, misperceiving a higher signal correlation in the stage game causes the agent to infer a lower price elasticity, as suggested

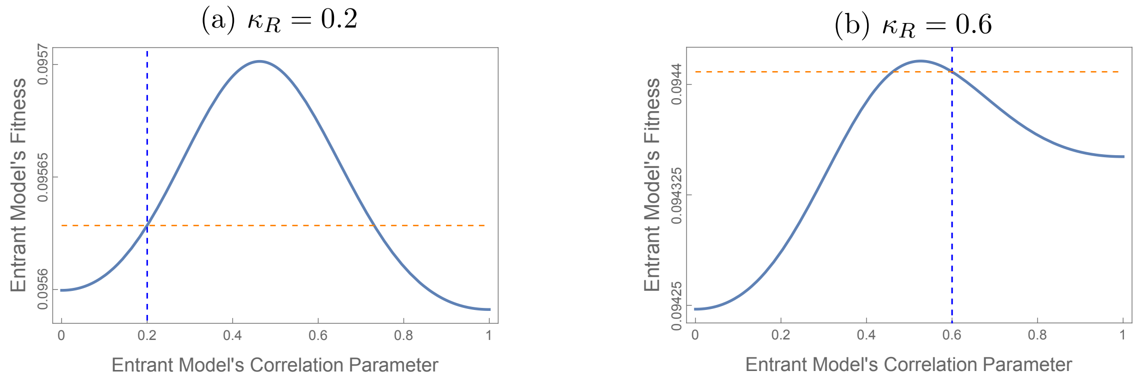

Figure 1: Fitness of flexible entrant models with different correlation perceptions, under uniform matching against residents with two different levels of κ_R. The true parameters are $\kappa^\bullet = 0.3$, $r^\bullet = 1$, $(\sigma_\zeta^\bullet)^2 = \sigma_\omega^2 = \sigma_\epsilon^2 = 1$. The dashed vertical line marks the resident's correlation parameter; the dashed horizontal line marks the resident's fitness.

by the intuition in the introduction. This intuition will drive the interaction between signal correlation misperception and misinference of the persistent price elasticity parameter in our main results in the next section.

4 Selecting Biases and the Role of the Learning Channel

We now turn to the selection of correlation perceptions and ask how the answer depends on whether agents have flexible models or dogmatic models. Throughout, we assume the true price shock variance exceeds the threshold L from Lemma 3. We first consider uniform matching with flexible models.

Proposition 1 (Uniform Matching Selects Projection Bias)**.** *Fix any $r^\bullet > 0$, $\kappa^\bullet \in [0,1]$ and $(\sigma_\zeta^\bullet)^2 \geq L$. Take any $\tilde{\kappa} \leq \kappa^\bullet$, and assume all agents have flexible models. Then there exists $\underline{\kappa} \in [0, \tilde{\kappa})$ and $\bar{\kappa} \in (\tilde{\kappa}, 1]$ so that whenever $(\kappa_R, \kappa_E) = (\tilde{\kappa}, \kappa)$ for $\kappa \in [\underline{\kappa}, \bar{\kappa}]$, there is a unique linear equilibrium with uniform matching.*

Furthermore, $\tilde{\kappa}$ is susceptible to invasion with uniform matching if $\kappa > \tilde{\kappa}$ and resistant to invasion with uniform matching if $\kappa < \tilde{\kappa}$.

A special case of the result is that $\kappa_R = \kappa^\bullet$, so that we find correctly specified residents are susceptible to invasion by projection-biased entrants under uniform matching. The intuition for this result follows from the observation that 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 (Fershtman and Judd, 1987). Here, misspecification about signal correlation leads to misinference about elasticity, which causes the entrants to respond credibly to their opponents' play in an overly aggressive manner.[7] The rational residents back down and yield a larger share of the surplus.

However, projection bias is beneficial only in small measures, as excessive aggression can lead to overproduction past the point where such commitments are beneficial—in other words, where the strategic benefits of the misspecification are outweighed by the direct losses from suboptimal production. Figure 1a illustrates this non-monotonicity of fitness in κ_E. Here, we present a representative plot of the entrant's fitness given a resident with $\kappa_R = 0.2 \leq \kappa^\bullet = 0.3$. While small increases in κ_E above κ_R improve entrant fitness, entrants no longer outperform residents if κ_E is close to one. Thus, there is a limited range of entrant κ_E values such that the entrant obtains higher welfare than the resident (i.e., where the graph of entrant fitness is above the horizontal dashed line).

Proposition 1's conclusion that residents are outperformed by entrants who have slightly higher perception of signal correlation also applies when the residents are themselves misspecified about signal correlation, provided $\kappa_R \leq \kappa^\bullet$. The intuition outlined above hinges on the assumption that the residents are not excessively aggressive on their own. If the resident were to exhibit a significant amount of projection bias (and thus under-infer price elasticity by the same channel discussed above), then there may no longer be a gain to using an even more aggressive strategy. This is illustrated in Figure 1b, which shows a representative plot of the entrant's fitness when κ_R is large—in particular, larger than $\kappa^\bullet$ by a sufficient amount. Here, while there are values of κ_E which outperform the resident, these values are *below* κ_R, rather than above it. In this case, using a κ_E closer to $\kappa^\bullet$ (and in particular, below κ_R) can yield higher fitness. This calculation shows that Proposition 1 cannot be strengthened to allow for arbitrary κ_R.

By contrast, assortative matching favors biases that lead to more *cooperative* behavior, and thus the commitment to aggression is detrimental to fitness. Correspondingly, we obtain the opposite result.

Proposition 2 (Assortative Matching Selects Correlation Neglect)**.** *Fix* $r^\bullet > 0$ *and* $(\sigma_\zeta^\bullet)^2 \geq L$. *Assume all agents have flexible models. Then,* κ_R *is susceptible to invasion with assortative*

[7] In Fershtman and Judd (1987), firms can pay managers a convex combination of profit and sales, with the main result being that a weight less than 1 should be placed on profit. Here, we show that similar commitments can emerge with misspecified signal correlation.

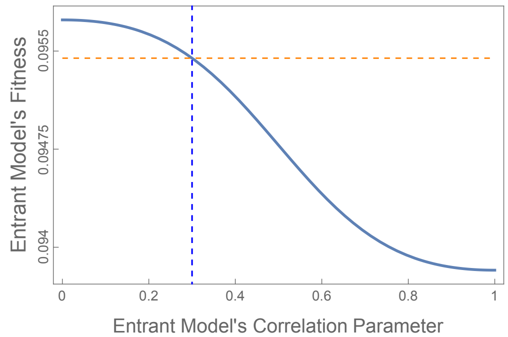

Figure 2: Fitness of flexible entrant models with different correlation perceptions, under assortative matching with a correctly specified resident model $\kappa_R = 0.3$. The true parameters are $\kappa^\bullet = 0.3$, $r^\bullet = 1$, $(\sigma_\zeta^\bullet)^2 = \sigma_\omega^2 = \sigma_\epsilon^2 = 1$. The dashed vertical line marks the objectively true correlation parameter; the dashed horizontal line marks the resident's fitness.

matching if $\kappa_E < \kappa_R$, and it is resistant to invasion with assortative matching if $\kappa_E > \kappa_R$.

Correlation neglect leads agents with flexible models to over-infer elasticity, enabling commitment to less aggressive behavior. Figure 2 shows a representative plot of entrant fitness under assortative matching for different correlation perceptions κ_E. Let α^{TEAM} denote the symmetric linear strategy profile that maximizes the sum of the two players' expected objective payoffs in the stage game. The proof of Proposition 2 shows that among symmetric strategy profiles, players' payoffs strictly decrease in aggressiveness in the region $\alpha > \alpha^{TEAM}$. For assortative matching and any $\kappa \in [0, 1]$, the linear equilibrium behavior in a group with correlation perception κ strictly increases in aggressiveness as κ grows, and this equilibrium play is always strictly more aggressive than α^{TEAM}.

Interestingly, in contrast to the case of uniform matching, with assortative matching there are no qualifiers regarding whether the misspecification induces excessive cooperation—more is always better. Lowering the perception of κ *always* confers an evolutionary advantage by bringing equilibrium play closer to α^{TEAM}—importantly, never reaching or exceeding it. Notice that the payoff of an entrant under assortative matching does not depend on the resident's model. The idea that assortativity leads to cooperation is similar to the findings of Alger and Weibull (2013); our contribution is to identify a particular channel for this to arise—showing that decreasing perceived signal correlation uniformly favors cooperation.

Our results thus far have focused on the selection of correlation misperceptions under flexible models. But as mentioned, the direct effect and the indirect effect of correlation misperception go in opposite directions. The next result shows that if agents have dogmatic models so that the learning channel is shut down, the conclusions of Propositions 1 and 2 can be reversed:

Proposition 3. *Let $\kappa^\bullet \in [0,1]$, $r^\bullet > 0$, $(\sigma_\zeta^\bullet)^2 \geq L$ be given and suppose all agents have dogmatic models whose fixed parameters about price elasticity and price shock variance are correct: $(\tilde{r}, \tilde{\sigma}_\zeta^2) = (r^\bullet, (\sigma_\zeta^\bullet)^2)$. Suppose $\kappa_R = \kappa^\bullet$.*

- *For any $\kappa_h > \kappa^\bullet$, κ_R is resistant to invasion from entrants with $\kappa_E = \kappa_h$ under uniform matching.*
- *For any $\kappa_l < \kappa^\bullet$, κ_R is resistant to invasion from entrants with $\kappa_E = \kappa_l$ under assortative matching.*

Proposition 3 shuts down the learning channel by considering agents who know the true value of r. This result identifies the indirect effect of elasticity misinference as responsible for the evolutionary advantage conferred by the misperceptions of κ analyzed in Propositions 1 and 2. Intuitively, Proposition 3 comes from the fact that increasing perceived signal correlation on its own leads to *less* aggressive strategies, as can be seen from the subjective best replies presented in Lemma 2. This is because production quantities are strategic substitutes, so that players who overestimate how much the opponent's expected quantity varies conditional on their own signals will react by making their own strategy depend less on the same signals. This force turns out to be weaker than the indirect effect of the κ misperception on the inference of r, which has the opposite impact on strategy aggressiveness. However, the indirect effect is only present when agents use flexible models. Thus, we obtain a sharp illustration of our main message that whether an error in higher-order beliefs can persist in a rational society may depend on whether the biased agents are open-minded or dogmatic about the values of the persistent parameters in the game.

In fact, the proof of Proposition 3 shows that it also applies in some situations with misspecified residents. In particular, suppose residents have flexible models and they perceive correlation to be κ_R, where possibly $\kappa_R \neq \kappa^\bullet$. Suppose entrants have dogmatic models whose fixed parameters about price elasticity and price shock variance are equal to those inferred by the residents when they play against each other in linear partial equilibrium. If $\kappa_R \leq \kappa^\bullet$, residents are resistant to invasion from entrants with any $\kappa_E > \kappa_R$ under uniform matching. For any κ_R, residents are resistant to invasion from residents with any $\kappa_E < \kappa_R$ under assortative matching. In other words, starting in the equilibrium of a society where everyone has the same (possibly misspecified) model, a given mutation in the correlation perception of a small group of agents leads to two opposite welfare consequences for them, depending on whether they can update their former belief about the persistent game parameters.

5 Extensions

5.1 Misperception of Others' Strategies

A key assumption of our framework, shared with the broader literature on the evolution of preferences, is that agents correctly know others' strategies in equilibrium. Here we discuss how this assumption shapes our conclusions on the selection of biases and the role of the learning channel.

We briefly describe how to modify our solution concept to accommodate misperceptions of others' strategies. For simplicity, we restrict attention to the case of uniform matching and correctly specified residents (who correctly know others' strategies in equilibrium). We suppose that, in addition to misperceiving κ, entrants also dogmatically misperceive the residents' strategy to be $\hat{\alpha}_R$, which may be different from their actual equilibrium strategy $\alpha_R^\bullet$. We consider both the case where the entrants have flexible models where price elasticity $\tilde{r}$ and price shock variance $\tilde{\sigma}_\zeta^2$ are free parameters, and the case where they have dogmatic models where price elasticity and price shock variance are fixed parameters at their true values $r^\bullet$ and $(\sigma_\zeta^\bullet)^2$.

In this modified setting, for entrants who have the misperceptions κ_E and $\hat{\alpha}_R$, a *linear equilibrium with strategy misperception* consists of strategies $\alpha_{R\to R}, \alpha_{R\to E}, \alpha_{E\to R}$ and beliefs $(\tilde{\kappa}_R, \tilde{r}_R, \tilde{\sigma}_{\zeta,R}^2)$, $(\kappa_E, \tilde{r}_E, \tilde{\sigma}_{\zeta,E}^2)$ such that:

- $\alpha_{R\to R}, \alpha_{R\to R}, (\tilde{\kappa}_R, \tilde{r}_R, \tilde{\sigma}_{\zeta,R}^2), (\tilde{\kappa}_R, \tilde{r}_R, \tilde{\sigma}_{\zeta,R}^2)$ are a linear partial equilibrium
- $\alpha_{E\to R}$ is an interim-stage best response against $\hat{\alpha}_R$, given the beliefs $(\kappa_E, \tilde{r}_R, \tilde{\sigma}_{\zeta,E}^2)$
- If entrants have flexible models, then $F^\bullet(\alpha_{E\to R}, \alpha_{R\to E}) = F_{\kappa_E, \tilde{r}_E, \tilde{\sigma}_{\zeta,E}^2}(\alpha_{E\to R}, \hat{\alpha}_R)$
- If entrants have dogmatic models, then $(\tilde{r}_E, \tilde{\sigma}_{\zeta,E}^2) = (r^\bullet, (\sigma_\zeta^\bullet)^2)$.

Here, the residents correctly know others' strategies. They draw inferences and choose subjectively optimal strategies just as in Definition 4. The entrants, given their beliefs about the stage-game parameters, choose a subjective best response not against the actual strategy of the residents, but against their misperception $\hat{\alpha}_R$ of the residents' strategy. Also, when entrants have flexible models, this strategy misperception also affects their inference of the persistent parameters, as they infer $\tilde{r}_E, \tilde{\sigma}^2_{\zeta,E}$ to rationalize the market price distribution under the hypothesis that residents use the strategy $\hat{\alpha}_R$.

By the same arguments as in the proof of Lemma 3, an entrant who repeatedly plays α_i against an opponent who plays $\alpha^\bullet_{-i}$, who misperceives signal correlation to be κ and misperceives the opponent's strategy to be $\hat{\alpha}_{-i}$, will infer price elasticity to be:

$$\hat{r} = r^\bullet \frac{\alpha_i + \alpha^\bullet_{-i}\psi(\kappa^\bullet)}{\alpha_i + \hat{\alpha}_{-i}\psi(\kappa)}.$$

This expression suggests that misperceptions of κ and misperceptions of others' strategies may have similar effects on inference and behavior. The next result verifies this intuition. When combined with a suitable strategy misperception, *every* correlation misperception is equally viable and the presence of the learning channel has no effect on the viability of a correlation misperception.

Proposition 4. *Fix any $r^\bullet > 0$, $\kappa^\bullet \in [0,1]$, $(\sigma^\bullet_\zeta)^2 \geq L$, and $\kappa_E \in [0,1]$. Let $U_i(\alpha_i, \alpha_{-i}; \kappa, r)$ and $\alpha_i^{BR}(\alpha_{-i}; \kappa, r)$ be as defined in Lemma 2, and define the Stackelberg payoff as*

$$\max_{\alpha_i} \quad U_i(\alpha_i, \alpha^{BR}_{-i}(\alpha_i; \kappa^\bullet, r^\bullet); \kappa^\bullet, r^\bullet). \tag{S}$$

For either entrants with flexible models or entrants with dogmatic models, there exists a strategy misperception $\hat{\alpha}_R$ such that in a society with uniform matching where residents are correctly specified and entrants have the misperceptions $(\kappa_E, \hat{\alpha}_R)$, there exists a unique linear equilibrium with strategy misperception. Furthermore, the entrants' fitness in this equilibrium is the Stackelberg payoff (S).

For any correlation misperception, there is some suitably complementary misperception of the residents' strategy to induce the entrants to play the Stackelberg strategy of the LQN game in equilibrium (taking into account the entrants' equilibrium misinference of the elasticity parameter in the case where they have flexible models). This gives them the Stackelberg payoff, which is the best they can hope for against rational residents. In particular, this shows that our conclusions about the comparative viability of different correlation misperceptions and about the role of the learning channel no longer apply in a setting where entrants can misperceive others' strategies. No correlation misperception is more viable than any other misperception, and the presence of the learning channel has no effect on the viability of any correlation misperception.

5.2 Evolutionary Advantage of Flexible Models with Correlation Misperception When There Are Multiple Environments

In this section, we provide another justification for the importance of the learning channel and the significance of higher-order errors. Roughly speaking, when agents must use the same model to learn in "multiple environments," flexible models with correlation misperception can perform better than any dogmatic model and better than first-order misperceptions of the game's persistent parameters. More precisely, Section 5.2.1 considers a setting where agents operate in multiple markets with different price elasticity parameters, so that the overall fitness of a model is given by a weighted sum of its equilibrium profits in the various markets. We show the resident rational model may be resistant to invasion from any dogmatic model but be susceptible to invasion from some flexible model. Section 5.2.2 compares the utility and inference consequences of different kinds of misperceptions in an alternative environment where the agents act as monopolists. We find that higher-order errors associated with correlation misperception cause no loss of profit and no belief distortion about price elasticity when agents are monopolists. By contrast, models that encode first-order errors about the slope or intercept of the demand curve always cause losses in the monopoly environment.

5.2.1 Stability and Invasion with Heterogeneous Markets

Suppose the agents interact in multiple markets with different true values of the price elasticity parameter. There are $M \geq 2$ different markets, indexed by $m = 1, 2, ..., M$, with M finite. The true price elasticity is $r^m > 0$ in market m and agents interact in this market with frequency $\phi^m > 0$ where $\sum_{m=1}^{M} \phi^m = 1$. Agents know which market they face in each period, but they are ex-ante uncertain about the values of the price elasticity parameters in different markets. Agents play a linear equilibrium in every market, so an agent can estimate different values of price elasticity for different markets (if r is a flexible parameter in their model).

Fix $(r^m), (\phi^m)$, the entrant model, and the resident model. We focus on the case of uniform matching. A *linear equilibrium with heterogeneous markets* consists of strategies $\alpha^m_{R\to R}, \alpha^m_{R\to E}, \alpha^m_{E\to R}$ and beliefs $(\tilde{\kappa}^m_R, \tilde{r}^m_R, (\tilde{\sigma}^2_{\zeta,R})^m), (\tilde{\kappa}^m_E, \tilde{r}^m_E, (\tilde{\sigma}^2_{\zeta,E})^m)$ for each market m, so that the strategies and beliefs in market m form a linear equilibrium with uniform matching under the true parameter r^m. We say that the resident model is *resistant to invasion* from the entrant model with uniform matching if

$$\sum_{m=1}^{M} \phi^m \cdot U^\bullet(\alpha^m_{R\to R}, \alpha^m_{R\to R}; r^m) \geq \sum_{m=1}^{M} \phi^m \cdot U^\bullet(\alpha^m_{E\to R}, \alpha^m_{R\to E}; r^m),$$

in every linear equilibrium with heterogeneous markets. We say it is susceptible to invasion if

$$\sum_{m=1}^{M} \phi^m \cdot U^\bullet(\alpha^m_{R\to R}, \alpha^m_{R\to R}; r^m) < \sum_{m=1}^{M} \phi^m \cdot U^\bullet(\alpha^m_{E\to R}, \alpha^m_{R\to E}; r^m),$$

in every linear equilibrium with heterogeneous markets. Here, $U^\bullet(\alpha_i, \alpha_{-i}; r^m)$ refers to the objective payoff when an agent uses α_i, their opponent uses α_{-i}, and the true price elasticity is r^m.

The next result shows that the rational resident model may be resistant to invasion from any dogmatic model, but be susceptible to invasion from a flexible entrant model with projection bias. The idea is that a dogmatic entrant model specifies the same fixed belief about price elasticity in all markets, which is beneficial for some values of the true price elasticity parameter but harmful for others. By contrast, a flexible entrant model can make different inferences about price elasticity in different markets and outperform the correctly specified resident model in every market.

Proposition 5. *Fix any $\kappa^\bullet \in [0,1]$, $(\sigma^\bullet_\zeta)^2 \geq L$. There exists a heterogeneous markets environment with $M = 2$ different markets such that the correctly specified resident model is resistant to invasion from any dogmatic entrant model. On the other hand, for any heterogeneous markets environment with any finite M, there exists a flexible model with projection bias so that the correctly specified resident model is susceptible to its invasion.*

5.2.2 Higher-Order Versus First-Order Misperceptions in Monopoly Markets

Now we consider an alternative stage game, which we call the *monopoly market*. Agents are paired to play this game and each chooses a quantity level. But, while the demand state $\omega \sim \mathcal{N}(0, \sigma^2_\omega)$ is commonly drawn for both players in the game, the two players do not interact strategically. Each player i is a monopolist and faces a market price P_i that does not depend on their opponent's quantity choice, with

$$P_i = \omega - r^\bullet q_i + \zeta_i$$

where price shock $\zeta_i \sim \mathcal{N}(0, (\sigma_\zeta^\bullet)^2)$ is drawn i.i.d. for the two players. As in the baseline stage game, each i observes a signal $s_i = \omega + \epsilon_i$ before choosing their quantity, where the joint distribution between ϵ_1 and ϵ_2 is the same as before. At the end of the match, agent i observes their signal s_i, their quantity choice q_i, and their market price P_i.

In addition to the misperception of signal correlation studied so far (a higher-order error about the belief of the opponent), we also consider two types of first-order misperceptions of the persistent parameters. Agents think that the market price is generated by $P_i = \omega + \theta - r^\bullet q_i + \zeta_i$ for some $\theta \in \mathbb{R}$. So, they need to form beliefs about the correlation parameter κ, price elasticity r, price shock variance σ_ζ^2, and intercept θ. A model dogmatically specifies the values of some parameters (*fixed parameters*) and lets the agent flexibly estimate the values of the other parameters (*free parameters*) in their respective domains. The domains of r and σ_ζ^2 are $[0, \infty)$ while the domain of θ is $\mathbb{R}$.

We define the equilibrium concept for a society with the monopoly market stage game, which generalizes the linear equilibrium concept by allowing for non-linear equilibrium strategies and equilibrium inferences that do not fully explain the distribution of consequences. Since the two players in the game do not affect each other's payoffs and observations, the matching assortativity is irrelevant. For the same reason, equilibrium only needs to specify a single strategy for each population, not multiple strategies to be played against different types of opponents.

An *equilibrium* consists of strategies $Q_R, Q_E : \mathbb{R} \to \mathbb{R}$ and beliefs $(\tilde{\kappa}_R, \tilde{r}_R, \tilde{\sigma}^2_{\zeta,R}, \theta_R), (\tilde{\kappa}_E, \tilde{r}_E, \tilde{\sigma}^2_{\zeta,E}, \theta_E)$ such that:

- For each player i and signal realization s_i, $Q_i(s_i)$ maximizes i's subjective expected utility given belief $(\tilde{\kappa}_i, \tilde{r}_i, \tilde{\sigma}^2_{\zeta,i}, \theta_i)$.
- For each player i, beliefs about the fixed parameters are as specified by i's model.
- For each player i, beliefs about the free parameters minimize the Kullback–Leibler divergence of the true distribution of consequences from the subjective distribution of consequences under the parameters and the strategy Q_i.

The flexible models that we have studied in the baseline environment correspond to models that have a fixed (and correct) parameter θ and have a fixed (and possibly wrong) parameter κ. The next result says that these models, in equilibrium, generate the objectively optimal payoffs for any misperception κ. That is, the higher-order errors that we have been studying do not lower profits in monopoly markets.

Proposition 6. *For any model where θ is fixed and correct, κ is fixed (and possibly wrong), and r, σ_ζ^2 are free parameters, the equilibrium objective expected utility of the agents who use this model is the highest possible across all strategies. In equilibrium, agents with this model infer the correct r and σ_ζ^2.*

Now consider models that encode first-order misperceptions of the persistent game parameters. The next result implies that models with a dogmatically wrong r parameter or dogmatically wrong θ parameter lead to losses, regardless of whether the other parameters are fixed or free.

Proposition 7. *For any model where r is fixed and wrong or θ is fixed and wrong, the equilibrium objective expected utility is strictly lower than the highest possible across all strategies.*

Thus, an evolutionary advantage of the correlation misperception compared to first-order misperceptions of the intercept or slope of the demand curve is that the former does not cause loss of profit in monopoly markets.

One story for why the higher-order error is more likely to emerge or persist is that agents need to do well in both duopoly markets and monopoly markets, so the invading bias must lead to weakly higher payoffs than that of the rational residents in both kinds of markets and strictly higher payoff in at least one. Proposition 7 shows that the correlation misperceptions that strictly increase the entrant's payoffs in the baseline duopoly markets perform just as well as the rational resident model in monopoly markets. While first-order misperceptions of the game parameters can also improve payoffs in duopoly markets, Proposition 7 implies they always lead to strict profit losses in monopoly markets.

5.3 More General Stage Games and Information Structures

We turn to general incomplete-information games and provide a condition for a model to be susceptible to invasion from a "nearby" misspecified model. This condition shows how assortativity and the learning channel shape the evolutionary selection of models for a broader class of stage games and biases.

Consider a stage game where a state of the world ω is drawn each time the game is played. 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 $\pi : \mathbb{Y} \to \mathbb{R}$ 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 models $(\Theta(\kappa))_{\kappa \in [\underline{\kappa}, \bar{\kappa}]}$. Let $\lambda = 0$ denote the case where matching is uniform and $\lambda = 1$ denote the case where it is assortative. Suppose the resident model in the society is $\Theta_R = \Theta(\kappa^{\bullet})$ and the entrant model is $\Theta_E = \Theta(\kappa)$ for some $\kappa \in [\underline{\kappa}, \bar{\kappa}]$. Each model Θ can contain both fixed parameters that dogmatically determine the agent's belief in equilibrium and free parameters whose values must be inferred to match the observable distribution over consequences given the equilibrium strategies.

As in Definition 4, a *linear equilibrium* for a given interaction structure $\lambda \in \{0, 1\}$ consists of each group g's beliefs μ_g about the persistent game parameters and each group g's strategy when playing each opponent group g' (denoted by $\sigma_{g \to g'} : S \to \mathbb{R}$) such that:

- Every agent maximizes their subjective expected utility in every match (given opponent's strategy and their belief about the persistent game parameters),
- Each group g's beliefs about the fixed parameters are as specified by g's model, and beliefs about the free parameters are self-confirming inferences — that is, parameter values that generate the observed distribution of consequences in the matches that group g faces with probability 1, and
- All strategies are linear functions of private signal realizations – that is, for every group g and every opponent group g', the strategy $\sigma_{g \to g'}(s_i) = \alpha_{g \to g'} \cdot s_i$ for some real number $\alpha_{g \to g'}$.

The next assumption requires there to be a unique linear equilibrium for either interaction structure and for any $\kappa \in [\underline{\kappa}, \bar{\kappa}]$. A significant portion of our analysis for our duopoly model aimed to demonstrate these existence and uniqueness properties. On the other hand, linear equilibria exist and are unique in a large class of games outside of the duopoly framework, and in particular in more general linear-quadratic-normal games under some conditions on the payoff functions (see, e.g., Angeletos and Pavan (2007)), and hence we expect this condition will also hold in broader environments.

Assumption 1. *Suppose there is a unique linear equilibrium for either $\lambda \in \{0,1\}$, with $\Theta_R = \Theta(\kappa^\bullet)$, $\Theta_E = \Theta(\kappa)$ for every $\kappa \in [\underline{\kappa}, \bar{\kappa}]$. Suppose the κ-indexed linear equilibrium strategies coefficients $\alpha_{R\to R}(\kappa), \alpha_{R\to E}(\kappa), \alpha_{E\to R}(\kappa), \alpha_{R\to R}(\kappa)$ are differentiable in κ. Finally, suppose that in the linear equilibrium with $\kappa = \kappa^\bullet$, $\alpha_{R\to R}(\kappa^\bullet)$ is objectively interim-optimal against itself.*[8]

Proposition 8. *Fix $\lambda \in \{0,1\}$ and let $\alpha^\bullet := \alpha_{R\to R}(\kappa^\bullet)$. Then, under Assumption 1, if*

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

then there exists some $\epsilon > 0$ so that $\Theta(\kappa^\bullet)$ is susceptible to invasion from models $\Theta(\kappa)$ with $\kappa \in (\kappa^\bullet, \kappa^\bullet + \epsilon] \cap [\underline{\kappa}, \bar{\kappa}]$, with λ interaction structure. Also, if

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

then there exists some $\epsilon > 0$ so that $\Theta(\kappa^\bullet)$ is susceptible to invasion from models $\Theta(\kappa)$ with $\kappa \in [\kappa^\bullet - \epsilon, \kappa^\bullet) \cap [\underline{\kappa}, \bar{\kappa}]$, with λ interaction structure. 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 model is susceptible to invasion from a "nearby" misspecified entrant model, as indexed by κ. The condition asks if a slight change in the entrant model's κ leads entrants' opponents to change their equilibrium actions such that the entrants become better off on average. These opponents are the residents under uniform matching $\lambda = 0$, so $\alpha'_{R\to E}(\kappa^\bullet)$ is relevant. These opponents are other entrants under assortative matching $\lambda = 1$, so $\alpha'_{E\to E}(\kappa^\bullet)$ is relevant.

Proposition 8 implies that one should only expect the correctly specified model to be resistant to invasion from all nearby models 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

[8]We say $\alpha_{R\to R}$ is objectively interim-optimal against itself if, for every $s_i \in S$, $\alpha_{R\to R}(\kappa^\bullet) \cdot s_i$ maximizes the agent's objective expected utility across all of $\mathbb{R}$ when $-i$ uses the same linear strategy $\alpha_{R\to R}(\kappa^\bullet)$.

agents face a decision problem where player 2's action does not affect player 1's payoffs, that is $\frac{\partial u_1^\bullet}{\partial q_2} = 0$. This condition sets the expectation to zero, so Proposition 8 never implies that the correctly specified model is susceptible to invasion from a misspecified model 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

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

Focusing on the case where all agents have flexible models, the proof of Proposition 1 shows that when $\lambda = 0$, $\alpha'_{R\to E}(\kappa^\bullet) < 0$. The proof of Proposition 2 shows that when $\lambda = 1$, $\alpha'_{E\to E}(\kappa^\bullet) > 0$. The uniqueness of linear equilibrium also follows from these results, for an open interval of κ containing $\kappa^\bullet$. Lemma 2 implies linear strategies played by two correctly specified players against each other are objectively interim-optimal. So, the conditions of Proposition 8 hold for $\lambda \in \{0, 1\}$, and we deduce that with flexible models, the correctly specified model is susceptible to invasion from slightly higher κ (for $\lambda = 0$) and slightly lower κ (for $\lambda = 1$).

6 Conclusion

The main message of this paper is that whether an error in higher-order belief in the stage game is likely to survive can depend on whether agents have dogmatic or flexible views about other persistent parameters of the stage game. In the context of an incomplete-information duopoly game, we show that the welfare implications of a higher-order misperception of the transient demand state depend crucially on whether people know the persistent price elasticity with certainty or estimate this elasticity from past prices. Working in a canonical linear-quadratic-normal game setting, we view our paper as illustrating the practical value of the evolutionary framework in terms of guiding our thinking about the viability of biases. More broadly, our results point out that the viability of a given error must be evaluated in the context of other factors, such as whether agents engage in inference about the persistent stage-game parameters. It may be worthwhile to investigate other factors that can enhance or hinder the viability of certain behavioral biases in future work.

Appendix

A Proofs

A.1 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 κ, 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 $\kappa \to 1$, 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 $\frac{1/\sigma_\epsilon^2}{1/\sigma_\epsilon^2 + 1/\sigma_\omega^2}$. □

A.2 Proof of Lemma 2

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

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

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 $\mathbb{E}[s_i^2] \cdot \left(\alpha_i\gamma - \frac{1}{2}r\alpha_i^2 - \frac{1}{2}r\psi(\kappa)\alpha_i\alpha_{-i} - \frac{1}{2}\alpha_i^2\right)$. 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 $-\frac{1}{2}r - \frac{1}{2} < 0$.

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

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

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

A.3 Proof of Lemma 3

Proof. Conditional on the signal s_i, the distribution of market price under the model $F_{\kappa,\hat{r},\hat{\sigma}_\zeta^2}$ 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 $F_{\kappa^\bullet,r^\bullet,(\sigma_\zeta^\bullet)^2}$ 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 for any α_i, α_{-i}. On the other hand, for any other value of $\hat{r}$, these means will not match for any $s_i \neq 0$.

Conditional on the signal s_i, the variance of market price under $F_{\kappa, r^\bullet \frac{\alpha_i+\alpha_{-i}\psi(\kappa^\bullet)}{\alpha_i+\alpha_{-i}\psi(\kappa)}, \hat{\sigma}_\zeta^2}$ is

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

By properties of the multivariate normal distribution, this conditional variance is constant in s_i. Let $L = \max_{\kappa\in[0,1], 0\le\alpha_i,\alpha_{-i}\le\gamma} \mathrm{Var}_\kappa \left[\omega - \frac{1}{2} r^\bullet \frac{\alpha_i+\alpha_{-i}\psi(\kappa^\bullet)}{\alpha_i+\alpha_{-i}\psi(\kappa)} \alpha_{-i} s_{-i} \mid s_i\right]$. This maximum exists and is finite since the expression is a continuous function of $\kappa, \alpha_i, \alpha_{-i}$ on the compact domain $[0,1] \times [0,\gamma]^2$. The conditional variance of market price under $F_{\kappa, r^\bullet \frac{\alpha_i+\alpha_{-i}\psi(\kappa^\bullet)}{\alpha_i+\alpha_{-i}\psi(\kappa)}, \hat{\sigma}_\zeta^2}$ is bounded by $L + \hat{\sigma}_\zeta^2$ whenever $0 \le \alpha_i, \alpha_{-i} \le \gamma$.

On the other hand, the variance of market price under $F_{\kappa^\bullet, r^\bullet, \sigma_\zeta^\bullet}$ is at least $(\sigma_\zeta^\bullet)^2$. Thus, whenever $(\sigma_\zeta^\bullet)^2 \ge L$, there exists a unique value of $\hat{\sigma}_\zeta^2$ such that the conditional variance under $F_{\kappa, r^\bullet \frac{\alpha_i+\alpha_{-i}\psi(\kappa^\bullet)}{\alpha_i+\alpha_{-i}\psi(\kappa)}, \hat{\sigma}_\zeta^2}$ is the same as that under $F_{\kappa^\bullet, r^\bullet, (\sigma_\zeta^\bullet)^2}$ given every s_i. □

A.4 Proof of Proposition 1

Proof. Take L as in Lemma 3. For $\kappa_R = \tilde{\kappa}$ and any $\kappa_E \in [0,1]$, consider a candidate linear equilibrium with strategies $0 \le \alpha_{R\to R}, \alpha_{E\to R}, \alpha_{R\to E} \le \gamma$, together with the self-confirming inferences given these strategies – such inferences exist and are unique by Lemma 3. From Lemma 3, the residents' belief must be $r_R := r^\bullet \frac{1+\psi(\kappa^\bullet)}{1+\psi(\kappa_R)}$. We note that $r^\bullet \le r_R \le 2r^\bullet$ since $\psi(\kappa_R) \in (0, \psi(\kappa^\bullet)]$.

Using the equilibrium belief of the resident, we must have $\alpha_{R\to R} = \alpha_i^{BR}(\alpha_{R\to R}; \kappa_R, r_R)$, so using the formula from Lemma 2 we find the unique solution $\alpha_{R\to R} = \frac{\gamma}{1+r_R+\frac{1}{2}r_R\psi(\kappa_R)}$. Next, we turn to $\alpha_{R\to E}, \alpha_{E\to R}$, and r_E, the entrant's self-confirming inference. For agents in each group to best respond to each others' play and for the entrant's inferences to be self-confirming, we must have $\alpha_{R\to E} = \frac{\gamma - \frac{1}{2}r_R\psi(\kappa_R)\alpha_{E\to R}}{1+r_R}$, $\alpha_{E\to R} = \frac{\gamma-\frac{1}{2}r_E\psi(\kappa_E)\alpha_{R\to E}}{1+r_E}$, and $r_E = r^\bullet \frac{\alpha_{E\to R}+\alpha_{R\to E}\psi(\kappa^\bullet)}{\alpha_{E\to R}+\alpha_{R\to E}\psi(\kappa_E)}$ from Lemma 3. We may rearrange the expression for $\alpha_{E\to R}$ to say $\alpha_{E\to R} = \gamma - r_E\alpha_{E\to R} - \frac{1}{2}r_E\psi(\kappa_E)\alpha_{R\to E}$. Substituting the expression of r_E into this expression of $\alpha_{E\to R}$, we get

$$\begin{aligned}\alpha_{E\to R} &= \gamma - r_E \cdot (\alpha_{E\to R} + \alpha_{R\to E}\psi(\kappa_E) - \frac{1}{2}\alpha_{R\to E}\psi(\kappa_E)) \\ &= \gamma - \frac{r^\bullet \alpha_{E\to R} + r^\bullet \alpha_{R\to E}\psi(\kappa^\bullet)}{\alpha_{E\to R} + \alpha_{R\to E}\psi(\kappa_E)} \cdot (\alpha_{E\to R} + \alpha_{R\to E}\psi(\kappa_E) - \frac{1}{2}\alpha_{R\to E}\psi(\kappa_E)) \\ &= \gamma - r^\bullet \alpha_{E\to R} - r^\bullet \alpha_{R\to E}\psi(\kappa^\bullet) + \frac{1}{2}\psi(\kappa_E)\alpha_{R\to E}\frac{r^\bullet \alpha_{E\to R} + r^\bullet \alpha_{R\to E}\psi(\kappa^\bullet)}{\alpha_{E\to R} + \alpha_{R\to E}\psi(\kappa_E)}\end{aligned}$$

Multiply by $\alpha_{E\to R} + \alpha_{R\to E}\psi(\kappa_E)$ on both sides and collect terms,

$$\begin{aligned}(\alpha_{E\to R})^2 \cdot [-1 - r^\bullet] + (\alpha_{E\to R}\alpha_{R\to E}) \cdot [-\psi(\kappa_E) - \frac{1}{2}r^\bullet\psi(\kappa_E) - r^\bullet\psi(\kappa^\bullet)] \\ - (\alpha_{R\to E})^2 \cdot [\frac{1}{2}r^\bullet\psi(\kappa^\bullet)\psi(\kappa_E)] + \gamma[\alpha_{E\to R} + \alpha_{R\to E}\psi(\kappa_E)] = 0. \quad (1)\end{aligned}$$

Consider the following quadratic function in x,

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

where $\ell(x) := \frac{\gamma - \frac{1}{2}r_R\psi(\kappa_R)x}{1+r_R}$ is a linear function in x. In a linear equilibrium, $\alpha_{E\to R}$ is a root of $H(x)$ in $[0, \frac{\gamma}{\frac{1}{2}r_R\psi(\kappa_R)}]$. To see why, if we were to have $\alpha_{E\to R} > \frac{\gamma}{\frac{1}{2}r_R\psi(\kappa_R)}$, then $\alpha_{R\to E} = 0$. In that case, $r_E = r^\bullet$ and so $\alpha_{E\to R} = \alpha_i^{BR}(0; \kappa_E, r^\bullet) = \frac{\gamma}{1+r^\bullet}$. Yet,

$$\frac{\gamma}{\frac{1}{2}r_R\psi(\kappa_R)} = \frac{\gamma}{\frac{1}{2}r^\bullet\frac{1+\psi(\kappa^\bullet)}{1+\psi(\kappa_R)}\psi(\kappa_R)} \geq \frac{\gamma}{\frac{1}{2}r^\bullet \cdot 2} = \frac{\gamma}{r^\bullet} \geq \frac{\gamma}{1+r^\bullet},$$

which contradicts $\alpha_{E\to R} > \frac{\gamma}{\frac{1}{2}r_R\psi(\kappa_R)}$. Conversely, for any root x^* of $H(x)$ in $[0, \frac{\gamma}{\frac{1}{2}r_R\psi(\kappa_R)}]$, there is a linear equilibrium where $\alpha_{E\to R} = x^*$, $\alpha_{R\to E} = \ell(x^*) \in [0, \gamma]$, and $r_E = r^\bullet\frac{\alpha_{E\to R}+\alpha_{R\to E}\psi(\kappa^\bullet)}{\alpha_{E\to R}+\alpha_{R\to E}\psi(\kappa_E)}$.

We now state and prove a useful claim.

Claim A.1. There exist some $\underline{\kappa}_1 < \kappa_R < \bar{\kappa}_1$ so that H has a unique root in $[0, \frac{\gamma}{\frac{1}{2}r_R\psi(\kappa_R)}]$ for all $\kappa_E \in [\underline{\kappa}_1, \bar{\kappa}_1] \cap [0, 1]$.

Proof. We show that when $\kappa_E = \kappa_R$, $H(x)$ (i) has a unique root in $[0, \frac{\gamma}{\frac{1}{2}r_R\psi(\kappa_R)}]$; (ii) $H(0) > 0$ and $H(\frac{\gamma}{\frac{1}{2}r_R\psi(\kappa_R)}) < 0$. By these two statements, since $H(x)$ is a continuous function of κ_E, there must exist some $\underline{\kappa}_1 < \kappa_R < \bar{\kappa}_1$ so that it continues to have a unique root in $[0, \frac{\gamma}{\frac{1}{2}r_R\psi(\kappa_R)}]$ for all $\kappa \in [\underline{\kappa}_1, \bar{\kappa}_1] \cap [0, 1]$.

Statement (i) has to do with the fact that when $\kappa_E = \kappa_R$, we need $\alpha_{E\to R} = \frac{\gamma - \frac{1}{2}r_R\psi(\kappa_R)\alpha_{R\to E}}{1+r_R}$ and $\alpha_{R\to E} = \frac{\gamma - \frac{1}{2}r_R\psi(\kappa_R)\alpha_{E\to R}}{1+r_R}$. These are linear best response functions with a slope of $-\frac{1}{2}\frac{r_R}{1+r_R}\psi(\kappa_R)$, which falls in $(-\frac{1}{2}, 0)$. So there can only be one solution to H in that region (even when we allow $\alpha_{E\to R} \neq \alpha_{R\to E}$), which is the symmetric equilibrium found before $\alpha_{E\to R} = \alpha_{R\to E} = \frac{\gamma}{1+r_R+\frac{1}{2}r_R\psi(\kappa_R)}$.

For Statement (ii), we evaluate $H(0) = -(\frac{\gamma}{1+r_R})^2\frac{1}{2}r^\bullet\psi(\kappa^\bullet)\psi(\kappa_R) + \frac{\gamma^2\psi(\kappa_R)}{1+r_R} = \frac{\psi(\kappa_R)\gamma^2}{1+r_R}(1 - \frac{(1/2)r^\bullet\psi(\kappa^\bullet)}{1+r_R}) > 0$ because $1 + r_R > (1/2)r^\bullet\psi(\kappa^\bullet)$, as $r_R \geq r^\bullet$. Then, we evaluate $H(\frac{\gamma}{\frac{1}{2}r_R\psi(\kappa_R)}) = (\frac{\gamma}{\frac{1}{2}r_R\psi(\kappa_R)})^2(-1 - r^\bullet) + \gamma\frac{\gamma}{\frac{1}{2}r_R\psi(\kappa_R)} = \frac{\gamma^2}{\frac{1}{2}r_R\psi(\kappa_R)}(1 - \frac{1+r^\bullet}{\frac{1}{2}r_R\psi(\kappa_R)})$. This is strictly negative because $r_R \leq 2r^\bullet$. $\square$

Returning to the proof of Proposition 1: by Claim A.1, for $\kappa_E \in [\underline{\kappa}_1, \bar{\kappa}_1] \cap [0, 1]$, entrants has only one possible belief about elasticity (denoted by $r_E(\kappa_E)$ in linear equilibrium), since there is only one possible outcome in the match between the entrants and the residents. So for every $\kappa_E \in [\underline{\kappa}_1, \bar{\kappa}_1] \cap [0, 1]$, there is a unique linear equilibrium, where equilibrium behavior is given as a function of κ_E by $\alpha(\kappa_E) = (\alpha_{R\to R}(\kappa_E), \alpha_{R\to E}(\kappa_E), \alpha_{E\to R}(\kappa_E))$.

Recall from Lemma 2 that the objective expected utility from playing α_i against an opponent who plays α_{-i} is $U^\bullet(\alpha_i, \alpha_{-i}) = \mathbb{E}[s_i^2] \cdot (\alpha_i\gamma - \frac{1}{2}r^\bullet\alpha_i^2 - \frac{1}{2}r^\bullet\psi(\kappa^\bullet)\alpha_i\alpha_{-i} - \frac{1}{2}\alpha_i^2)$. If $-i$ plays the best response under beliefs (κ_R, r_R), then the objective expected utility of choosing α_i is $\bar{U}_i(\alpha_i) := \mathbb{E}[s_i^2] \cdot (\alpha_i\gamma - \frac{1}{2}r^\bullet\alpha_i^2 - \frac{1}{2}r^\bullet\psi(\kappa^\bullet)\alpha_i\frac{\gamma - \frac{1}{2}r_R\psi(\kappa_R)\alpha_i}{1+r_R} - \frac{1}{2}\alpha_i^2)$. The derivative in α_i has the same sign as:

$$\bar{U}_i'(\alpha_i) \propto \gamma - r^\bullet\alpha_i - \frac{1}{2}r^\bullet\psi(\kappa^\bullet)[\frac{\gamma - \frac{1}{2}r_R\psi(\kappa_R)\alpha_i}{1+r_R} - \alpha_i\frac{\frac{1}{2}r_R\psi(\kappa_R)}{1+r_R}] - \alpha_i.$$

So, $\bar{U}_i'(\alpha_{RR})$ has the same sign as:

$$\begin{aligned} &\gamma - r^\bullet\alpha_{RR} - \frac{1}{2}r^\bullet\psi(\kappa^\bullet)[\frac{\gamma - \frac{1}{2}r_R\psi(\kappa_R)\alpha_{RR}}{1+r_R} - \alpha_{RR}\frac{\frac{1}{2}r_R\psi(\kappa_R)}{1+r_R}] - \alpha_{RR} \\ =&\gamma - r^\bullet\alpha_{RR} - \frac{1}{2}r^\bullet\psi(\kappa^\bullet)\frac{1}{1+r_R}[\gamma - r_R\psi(\kappa_R)\alpha_{RR}] - \alpha_{RR} \\ =&\gamma \cdot (1 - \frac{1}{2}r^\bullet\psi(\kappa^\bullet)\frac{1}{1+r_R}) + \alpha_{RR}[\frac{1}{2}r^\bullet\psi(\kappa^\bullet)\frac{r_R\psi(\kappa_R)}{1+r_R} - r^\bullet - 1]. \end{aligned}$$

Making the substitution $\alpha_{RR} = \frac{\gamma}{1+r_R+\frac{1}{2}r_R\psi(\kappa_R)}$, we get:

$$1 - \frac{1}{2}r^\bullet\psi(\kappa^\bullet)\frac{1}{1+r_R} + \frac{1}{1+r_R+\frac{1}{2}r_R\psi(\kappa_R)} \cdot [\frac{1}{2}r^\bullet\psi(\kappa^\bullet)\frac{r_R\psi(\kappa_R)}{1+r_R} - r^\bullet - 1].$$

Multiply through by $(1+r_R)(1+r_R+\frac{1}{2}r_R\psi(\kappa_R))$, we have

$$\begin{aligned}
&(1+r_R)(1+r_R+\frac{1}{2}r_R\psi(\kappa_R)) - \frac{1}{2}r^\bullet\psi(\kappa^\bullet)[1+r_R+\frac{1}{2}r_R\psi(\kappa_R)] + [\frac{1}{2}r^\bullet\psi(\kappa^\bullet)r_R\psi(\kappa_R)] - [1+r^\bullet]\cdot(1+r_R)\\
=&(1+r_R)(1+r_R+\frac{1}{2}r_R\psi(\kappa_R)) - \frac{1}{2}r^\bullet\psi(\kappa^\bullet)[1+r_R] + [\frac{1}{4}r^\bullet\psi(\kappa^\bullet)r_R\psi(\kappa_R)] - [1+r^\bullet]\cdot(1+r_R)\\
=&(1+r_R)(1+r_R+\frac{1}{2}r_R\psi(\kappa_R) - \frac{1}{2}r^\bullet\psi(\kappa^\bullet) - 1 - r^\bullet) + [\frac{1}{4}r^\bullet\psi(\kappa^\bullet)r_R\psi(\kappa_R)]\\
=&(1+r_R)[(r_R)(1+\frac{1}{2}\psi(\kappa_R)) - r^\bullet(1+\frac{1}{2}\psi(\kappa^\bullet))] + [\frac{1}{4}r^\bullet\psi(\kappa^\bullet)r_R\psi(\kappa_R)].
\end{aligned}$$

A sufficient condition for $\bar{U}_i'(\alpha_{RR}) > 0$ is for $\frac{(r_R)(1+\frac{1}{2}\psi(\kappa_R))}{r^\bullet(1+\frac{1}{2}\psi(\kappa^\bullet))} \geq 1$. We have:

$$\frac{(r_R)(1+\frac{1}{2}\psi(\kappa_R))}{r^\bullet(1+\frac{1}{2}\psi(\kappa^\bullet))} = \frac{1+\psi(\kappa^\bullet)}{1+\psi(\kappa_R)} \cdot \frac{(1+\frac{1}{2}\psi(\kappa_R))}{1+\frac{1}{2}\psi(\kappa^\bullet)} = \frac{1+\psi(\kappa^\bullet)}{1+\frac{1}{2}\psi(\kappa^\bullet)} \cdot \frac{1+\frac{1}{2}\psi(\kappa_R)}{1+\psi(\kappa_R)}$$

But note that this term is 1 when $\kappa_R = \kappa^\bullet$, and $\frac{d}{dx}[\frac{1+\frac{1}{2}\psi(x)}{1+\psi(x)}]$ has the same sign as $\frac{1}{2}\psi'(x)(1+\psi(x)) - (1+\frac{1}{2}\psi(x))\psi'(x) = \frac{1}{2}\psi'(x) - \psi'(x) < 0$ for all $x \in [0,1]$, since $\psi'(x) > 0$ for all $x \in [0,1]$. This tells us that in fact, for any $\kappa_R \leq \kappa^\bullet$, we get $\bar{U}_i'(\alpha_{RR}) > 0$.

Therefore, if we can show that $\alpha'_{E\to R}(\kappa_R) > 0$, then there exists some $\underline{\kappa}_1 \leq \underline{\kappa} < \kappa_R < \bar{\kappa} \leq \bar{\kappa}_1$ so that for every $\kappa_E \in [\underline{\kappa}, \bar{\kappa}] \cap [0,1]$, $\kappa_E \neq \kappa_R$ entrants have strictly higher or strictly lower equilibrium fitness in the unique linear equilibrium than residents, depending on the sign of $\kappa_E - \kappa_R$.

Consider again the quadratic function $H(x)$ in Equation (2) and implicitly characterize the unique root x in $[0, \frac{\gamma}{\frac{1}{2}r_R\psi(\kappa_R)}]$ as a function of κ_E in a neighborhood around κ_R. Denote this root by α^M, let $D := \frac{d\alpha^M}{d\psi(\kappa_E)}$ and also note $\frac{d\ell(\alpha^M)}{d\psi(\kappa_E)} = \frac{-r_R}{2(1+r_R)}\psi(\kappa_R)\cdot D$. We have

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

Evaluate at $\kappa_E = \kappa_R$, noting that $\alpha^M(\kappa_R) = \ell(\alpha^M(\kappa_R)) = x^* := \frac{\gamma}{1+r_R+\frac{1}{2}\psi(\kappa_R)r_R}$. The terms without D are:

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

Now we show $\frac{1+r^\bullet+\frac{1}{2}\psi(\kappa^\bullet)r^\bullet}{1+r_R+\frac{1}{2}\psi(\kappa_R)r_R} \leq 1$ for $\kappa_R \leq \kappa^\bullet$. We have $r_R = r^\bullet \frac{1+\psi(\kappa^\bullet)}{1+\psi(\kappa_R)}$, so $r_R + \frac{1}{2}\psi(\kappa_R)r_R = r^\bullet[1+\psi(\kappa^\bullet)] \cdot \frac{1+\frac{1}{2}\psi(\kappa_R)}{1+\psi(\kappa_R)}$. As shown before, $\frac{d}{dx}[\frac{1+\frac{1}{2}\psi(x)}{1+\psi(x)}] < 0$, so $1 + r_R + \frac{1}{2}\psi(\kappa_R)r_R \geq 1 + r^\bullet + \frac{1}{2}\psi(\kappa^\bullet)r^\bullet$ for $\kappa_R \leq \kappa^\bullet$. Thus, the terms without D are larger than $x^* \cdot \left[-\gamma + \frac{1}{2}x^* r^\bullet + \gamma\right]$, so they are strictly positive.

The coefficient in front of D is:

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

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

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

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

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

For the terms without $\psi(\kappa^\bullet)$ and $\psi(\kappa_R)$, we have $1 + r_R - 2 - 2r^\bullet < 0$, since $r_R \leq 2 \cdot r^\bullet$.

For the terms with the first power of $\psi(\kappa^\bullet)$ or $\psi(\kappa_R)$, the only positive term is $\frac{1}{2}\psi(\kappa_R)r_R$ and the two negative terms are $\psi(\kappa^\bullet)(-r^\bullet)$ and $\psi(\kappa_R)(-\frac{1}{2}r^\bullet - 1)$. We note that $\psi(\kappa^\bullet)(r^\bullet) \geq$

$\psi(\kappa_R)r_R$. This is because $\frac{\psi(\kappa_R)r_R}{\psi(\kappa^\bullet)(r^\bullet)} = \frac{\psi(\kappa_R)\cdot r^\bullet \cdot \frac{1+\psi(\kappa^\bullet)}{1+\psi(\kappa_R)}}{\psi(\kappa^\bullet)(r^\bullet)} = \frac{1+\psi(\kappa^\bullet)}{\psi(\kappa^\bullet)} \cdot \frac{\psi(\kappa_R)}{1+\psi(\kappa_R)}$. This ratio is 1 when $\kappa_R = \kappa^\bullet$, and $\frac{d}{dx}[\frac{\psi(x)}{1+\psi(x)}] \propto \psi'(x)(1+\psi(x)) - \psi(x)\psi'(x) = \psi'(x) > 0$. So, when $k_R \leq \kappa^\bullet$, we have $\frac{\psi(\kappa_R)r_R}{\psi(\kappa^\bullet)(r^\bullet)} \leq 1$ and hence $\psi(\kappa^\bullet)(r^\bullet) \geq \psi(\kappa_R)r_R$. So, the positive term $\frac{1}{2}\psi(\kappa_R)r_R$ is canceled out by half of $\psi(\kappa^\bullet)(-r^\bullet)$, and we still have $\frac{1}{2}\psi(\kappa^\bullet)(-r^\bullet)$ leftover. The negative term $\psi(\kappa_R)(-\frac{1}{2}r^\bullet - 1)$ also remains.

For the terms with the third power of $\psi(\kappa^\bullet)$ and $\psi(\kappa_R)$, the only positive term is $\psi(\kappa_R)^2 \cdot \psi(\kappa^\bullet)\frac{r^\bullet r_R}{2(1+r_R)}$. But this term is smaller than $\frac{1}{2}\psi(\kappa^\bullet)r^\bullet$, since $\psi(\kappa_R)^2 \cdot \frac{r_R}{1+r_R} < 1$. Therefore it is canceled out by the leftover term $\frac{1}{2}\psi(\kappa^\bullet)(-r^\bullet)$ from before.

Finally, for the terms with the second power $\psi(\kappa^\bullet)$ and $\psi(\kappa_R)$, we evaluate:

$$\frac{r_R}{2(1+r_R)}\psi(\kappa_R)[\psi(\kappa_R)+\frac{1}{2}r^\bullet\psi(\kappa_R)+r^\bullet\psi(\kappa^\bullet)]-\frac{\psi(\kappa_R)^2 r_R}{2(1+r_R)}\cdot(1+r_R) \leq \frac{r_R}{2(1+r_R)}\psi(\kappa_R)[r^\bullet\psi(\kappa^\bullet)],$$

using the fact that $r^\bullet \leq r_R$. Combining $\frac{r_R}{2(1+r_R)}\psi(\kappa_R)[r^\bullet\psi(\kappa^\bullet)]$ with the unused term $\psi(\kappa_R)(-\frac{1}{2}r^\bullet - 1)$ from the first power of $\psi(\kappa_R)$, the sum is negative since $\frac{r_R}{1+r_R}\psi(\kappa^\bullet) < 1$.

Thus the coefficient in front of D is strictly negative. This shows $D(\kappa_R) > 0$. 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. Hence, we get $\alpha'_{E\to R}(\kappa_R) > 0$ as desired. □

A.5 Proof of Proposition 2

Proof. We will show that in every linear equilibrium: (i) for each $g \in \{R, E\}$, the inferred elasticity under κ_g is $\frac{1+\psi(\kappa^\bullet)}{1+\psi(\kappa_g)}r^\bullet$; (ii) for each $g \in \{R, E\}$, $\alpha_{g\to g} = \frac{\gamma}{1+\frac{r^\bullet}{2}(1+\psi(\kappa^\bullet))+\frac{r^\bullet}{2}(\frac{1+\psi(\kappa^\bullet)}{1+\psi(\kappa_g)})}$; (iii) the equilibrium fitness of group g is weakly higher than that of group g' if and only if $\kappa_g \leq \kappa_{g'}$.

Take L as in Lemma 3. In any linear equilibrium, by Lemma 3, group g agents infer elasticity $r_i^{INF}(\alpha_{g\to g}, \alpha_{g\to g}; \kappa^\bullet, \kappa_g, r^\bullet) = \frac{\alpha_{g\to g}+\alpha_{g\to g}\psi(\kappa^\bullet)}{\alpha_{g\to g}+\alpha_{g\to g}\psi(\kappa_g)}r^\bullet = \frac{1+\psi(\kappa^\bullet)}{1+\psi(\kappa_g)}r^\bullet$, proving (i).

Given this belief, we must have $\alpha_{g\to g} = \frac{\gamma - \frac{1}{2}\frac{1+\psi(\kappa^\bullet)}{1+\psi(\kappa_g)}r^\bullet\psi(\kappa_g)\alpha_{g\to g}}{1+\frac{1+\psi(\kappa^\bullet)}{1+\psi(\kappa_g)}r^\bullet}$ by Lemma 2. Rearranging yields $\alpha_{g\to g} = \frac{\gamma}{1+\frac{r^\bullet}{2}(1+\psi(\kappa^\bullet))+\frac{r^\bullet}{2}(\frac{1+\psi(\kappa^\bullet)}{1+\psi(\kappa)})}$, 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 function is strictly concave and quadratic 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,

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

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

A.6 Proof of Proposition 3

Proof. Take L as in Lemma 3. Consider the more general case where residents have flexible models with the (mis)perception κ_R. Let r_R be their belief about the price elasticity in linear equilibrium, which is not affected by the entrants' behavior. In particular, if $\kappa_R = \kappa^\bullet$, then $r_R = r^\bullet$. Suppose the entrant has the dogmatic model with fixed parameters $\kappa_E, r_R, (\sigma_{\zeta,R})^2$ for some $\kappa_E \in [0, 1]$. Following the same steps of the proof of Proposition 1, there exists exactly one linear equilibrium, and it involves residents playing $\frac{\gamma}{1+r_R+\frac{1}{2}r_R\psi(\kappa_R)}$ against each other and believing price elasticity to be r_R.

For the first claim of the proposition, suppose $\kappa_R \le \kappa^\bullet$ and $\kappa_E \ge \kappa_R$. We begin by showing that the equilibrium strategy that the entrants use against residents is strictly decreasing in κ_E.

Consider a linear equilibrium where in the matches between entrants and residents, the entrants use α_E and the residents use α_R. By Lemma 2, the best response function of the residents and the entrants imply that

$$\alpha_R = \frac{\gamma - \frac{1}{2} r_R \psi(\kappa_R)\alpha_E}{1 + r_R}$$

and

$$\alpha_E = \frac{\gamma - \frac{1}{2} r_R \psi(\kappa_E)\alpha_R}{1 + r_R}.$$

Making the substitution $\alpha_R = \frac{\gamma - \frac{1}{2} r_R \psi(\kappa_R)\alpha_E}{1+r_R}$ in the expression for α_E, we find that the value of α_E is pinned down by

$$\alpha_E = \frac{\gamma - \frac{1}{2} r_R \psi(\kappa_E) \left[\frac{\gamma - \frac{1}{2} r_R \psi(\kappa_R)\alpha_E}{1+r_R} \right]}{1+r_R}.$$

Multiplying both sides by $(1+r_R)^2$ and rearranging, we get:

$$\alpha_E = \frac{\gamma(1 + r_R - \frac{1}{2} r_R \psi(\kappa_E))}{(1+r_R)^2 - \frac{1}{4}(r_R)^2 \psi(\kappa_E)\psi(\kappa_R)}.$$

So, $\frac{d\alpha_E}{d\kappa_E}$ has the same sign as:

$$-\frac{1}{2} r_R \gamma \psi'(\kappa_E) \left[(1+r_R)^2 - \frac{1}{4}(r_R)^2 \psi(\kappa_E)\psi(\kappa_R) \right] + \gamma(1 + r_R - \frac{1}{2} r_R \psi(\kappa_E)) \cdot \frac{1}{4}(r_R)^2 \psi'(\kappa_E)\psi(\kappa_R). \tag{3}$$

We note that

$$\begin{aligned} (1+r_R)^2 - \frac{1}{4}(r_R)^2 \psi(\kappa_E)\psi(\kappa_R) &\geq (1+r_R)^2 - \frac{1}{4}(1+r_R)^2 \psi(\kappa_E)\psi(\kappa_R) \\ &\geq \frac{3}{4}(1+r_R)^2 \end{aligned}$$

since $\psi(\kappa_E), \psi(\kappa_R) \leq 1$. Also, we have

$$\begin{aligned} \gamma(1 + r_R - \frac{1}{2} r_R \psi(\kappa_E)) \cdot \frac{1}{4}(r_R)^2 \psi'(\kappa_E)\psi(\kappa_R) &\leq \gamma(1+r_R) \cdot \frac{1}{4}(r_R)^2 \psi'(\kappa_E)\psi(\kappa_R) \\ &\leq \frac{1}{4}\gamma(1+r_R)^2 (r_R) \psi'(\kappa_E)\psi(\kappa_R) \\ &\leq \frac{1}{4}\gamma(1+r_R)^2 (r_R) \psi'(\kappa_E), \end{aligned}$$

again using the fact that $\psi(\kappa_R) \leq 1$. Therefore, the expression in Equation (3) is no larger than

$$-\frac{1}{2} r_R \gamma \psi'(\kappa_E) \cdot \frac{3}{4}(1+r_R)^2 + \frac{1}{4}\gamma(1+r_R)^2 (r_R)\psi'(\kappa_E) = \gamma \psi'(\kappa_E) r_R (1+r_R)^2 \cdot (\frac{-1}{8}) < 0,$$

since $\gamma, \psi'(\kappa_E), r_R$ are all strictly positive. Thus we conclude $\frac{d\alpha_E}{d\kappa_E} < 0$ for every $\kappa_E \in [0,1]$.

Next, consider the function $\bar{U}_i(\alpha_i)$ (for the case of residents having the correlation perception κ_R) defined in the proof of Proposition 1, which describes an agent's objective expected utility from playing α_i when their opponent plays a subjective best response to α_i under the beliefs (κ_R, r_R). The fitness of entrants with misperception κ_E is $\bar{U}_i(\alpha_E(\kappa_E))$, so combined with the fact that $\frac{d\alpha_E}{d\kappa_E} < 0$, to establish the first claim of the proposition we just need to show that $\bar{U}_i'(\alpha_i)$ is strictly positive for all $\alpha_i \leq \alpha_{RR}$, where α_{RR} is the equilibrium strategy that the residents use against each other. The arguments in the proof of Proposition 1 show that $\bar{U}_i'(\alpha_i)$ has the same sign as

$$\gamma \cdot (1 - \frac{1}{2}r^\bullet\psi(\kappa^\bullet)\frac{1}{1+r_R}) + \alpha_i \cdot [\frac{1}{2}r^\bullet\psi(\kappa^\bullet)\frac{r_R\psi(\kappa_R)}{1+r_R} - r^\bullet - 1].$$

This expression is linear in α_i. The proof of Proposition 1 shows that it is strictly positive when $\alpha_i = \alpha_{RR}$. When $\alpha_i = 0$, this expression is equal to $\gamma \cdot (1 - \frac{1}{2}r^\bullet\psi(\kappa^\bullet)\frac{1}{1+r_R})$, which is also strictly positive since $\frac{1}{2}r^\bullet\psi(\kappa^\bullet)\frac{1}{1+r_R} \leq 1/2$ (as we know from Lemma 3, $r_R \geq r^\bullet$ when $\kappa_R \leq \kappa^\bullet$). So this expression is always strictly positive for every $\alpha_i \in [0, \alpha_{RR}]$, which shows $\bar{U}_i'(\alpha_i) > 0$ for all $\alpha_i \leq \alpha_{RR}$.

Next, we turn to $\alpha_{E\to E}(\kappa)$ with assortative matching. Using the expression for α_i^{BR} in Lemma 2, we find that $\alpha_{E\to E}(\kappa) = \frac{\gamma}{1+r_R+\frac{1}{2}r_R\psi(\kappa)}$. Since $\psi' > 0$, we have $\alpha_{E\to E}(\kappa)$ is strictly larger than $\alpha_{R\to R} = \frac{\gamma}{1+r_R+\frac{1}{2}r_R\psi(\kappa_R)}$ when $\kappa < \kappa_R$. From the proof of Proposition 2, we know that objective payoffs in the stage game are strictly decreasing in linear strategies larger than the team solution $\alpha_{TEAM} = \frac{\gamma}{1+r^\bullet+r^\bullet\psi(\kappa^\bullet)}$. Since $\alpha_{E\to E}(\kappa) > \alpha_{R\to R} > \alpha_{TEAM}$, we conclude the entrants with $\kappa_E = \kappa_l$ have strictly lower fitness than residents with κ_R in the unique linear equilibrium with assortative matching for any $\kappa_l < \kappa_R$. This argument establishes the second claim. □

A.7 Proof of Proposition 4

Proof. Fix any arbitrary entrant misperception κ. We define α_i^{SL} and α_{-i}^{SF} to be the (objective) Stackelberg leader and follower strategies under the true parameters. That is, α_i^{SL} solves $\max_{\alpha_i} U_i(\alpha_i, \alpha_{-i}^{BR}(\alpha_i; \kappa^\bullet, r^\bullet); \kappa^\bullet, r^\bullet)$ where $\alpha_{-i}^{BR}(\alpha_i; \kappa^\bullet, r^\bullet)$ is the rational best response against α_i, and $\alpha_{-i}^{SF} = \alpha_{-i}^{BR}(\alpha_i^{SL}; \kappa^\bullet, r^\bullet)$. Using the expression for the α_{-i}^{BR} function from Lemma 2, we get:

$$\alpha_i^{SL} = \frac{\gamma(2(1+r^\bullet) - \psi(\kappa^\bullet)r^\bullet)}{2 + 2r^\bullet(2+r^\bullet) - \psi(\kappa^\bullet)^2(r^\bullet)^2}$$
$$\alpha_{-i}^{SF} = \frac{\gamma}{1+r^\bullet} - \frac{1}{2}\alpha_i^{SL}\psi(\kappa^\bullet)\frac{r^\bullet}{1+r^\bullet}.$$

Note that $\alpha_i^{SL} > 0$ since $\psi(\kappa^\bullet) < 1$; we also have $\alpha_{-i}^{SF} = \frac{\gamma}{2(1+r^\bullet)}\left(2 - \frac{(2(1+r^\bullet)-\psi(\kappa^\bullet)r^\bullet)\psi(\kappa^\bullet)r^\bullet}{2+2r^\bullet(2+r^\bullet)-\psi(\kappa^\bullet)^2(r^\bullet)^2}\right) > 0$. Following the same steps as in Lemma 3, when i uses α_i^{SL} and $-i$ uses α_{-i}^{SF}, then i with misperception κ and $\hat{\alpha}_{-i}$ misinfers:

$$\hat{r} = r^\bullet \frac{\alpha_i^{SL} + \alpha_{-i}^{SF}\psi(\kappa^\bullet)}{\alpha_i^{SL} + \hat{\alpha}_{-i}\psi(\kappa)}.$$

When paired with an appropriate misinference of $(\sigma_\zeta^\bullet)^2$, this misinference is self-confirming. We show that there is some $\hat{\alpha}_{-i}$ such that $\hat{r}$ induces the entrant to follow the Stackelberg strategy. Again by Lemma 2, we have:

$$\alpha_i^{BR}(\hat{\alpha}_{-i};\kappa;r^\bullet \frac{\alpha_i^{SL} + \alpha_{-i}^{SF}\psi(\kappa^\bullet)}{\alpha_i^{SL} + \hat{\alpha}_{-i}\psi(\kappa)}) = \frac{\gamma - \frac{1}{2}\psi(\kappa)r^\bullet(\alpha_i^{SL} + \alpha_{-i}^{SF}\psi(\kappa^\bullet))\frac{\hat{\alpha}_{-i}}{\alpha_i^{SL}+\hat{\alpha}_{-i}\psi(\kappa)}}{1 + r^\bullet\frac{\alpha_i^{SL}+\alpha_{-i}^{SF}\psi(\kappa^\bullet)}{\alpha_i^{SL}+\hat{\alpha}_{-i}\psi(\kappa)}}$$

We show that (1) as $\hat{\alpha}_{-i} \to \infty$, $\alpha_i^{BR}(\hat{\alpha}_{-i};\kappa;r^\bullet\frac{\alpha_i^{SL}+\alpha_{-i}^{SF}\psi(\kappa^\bullet)}{\alpha_i^{SL}+\hat{\alpha}_{-i}\psi(\kappa)}) > \alpha_i^{SL}$, and that (2) as $\hat{\alpha}_{-i} \to 0$, $\alpha_i^{BR}(\hat{\alpha}_{-i};\kappa;r^\bullet\frac{\alpha_i^{SL}+\alpha_{-i}^{SF}\psi(\kappa^\bullet)}{\alpha_i^{SL}+\hat{\alpha}_{-i}\psi(\kappa)}) < \alpha_i^{SL}$. Continuity of $\alpha_i^{BR}(\hat{\alpha}_{-i};\kappa;r^\bullet\frac{\alpha_i^{SL}+\alpha_{-i}^{SF}\psi(\kappa^\bullet)}{\alpha_i^{SL}+\hat{\alpha}_{-i}\psi(\kappa)})$ in $\hat{\alpha}_{-i}$ then completes the proof of the Proposition, since the intermediate value theorem implies some $\hat{\alpha}_{-i}$ such that $\alpha_i^{BR}(\hat{\alpha}_{-i};\kappa;r^\bullet\frac{\alpha_i^{SL}+\alpha_{-i}^{SF}\psi(\kappa^\bullet)}{\alpha_i^{SL}+\hat{\alpha}_{-i}\psi(\kappa)}) = \alpha_i^{SL}$. At this misperception of $\hat{\alpha}_{-i}$, there is a linear equilibrium where, in the matches between an entrant and a resident, the entrant uses α_i^{SL}, the resident uses α_{-i}^{SF}, and the entrant infers $\hat{r} = r^\bullet\frac{\alpha_i^{SL}+\alpha_{-i}^{SF}\psi(\kappa^\bullet)}{\alpha_i^{SL}+\hat{\alpha}_{-i}\psi(\kappa)}$ which rationalizes the entrants playing α_i^{SL}. Notice that:

$$\lim_{\hat{\alpha}_{-i}\to\infty} \alpha_i^{BR}(\hat{\alpha}_{-i};\kappa;r^\bullet\frac{\alpha_i^{SL} + \alpha_{-i}^{SF}\psi(\kappa^\bullet)}{\alpha_i^{SL} + \hat{\alpha}_{-i}\psi(\kappa)}) = \gamma - \frac{1}{2}r^\bullet(\alpha_i^{SL} + \alpha_{-i}^{SF}\psi(\kappa^\bullet))$$
$$\lim_{\hat{\alpha}_{-i}\to 0} \alpha_i^{BR}(\hat{\alpha}_{-i};\kappa;r^\bullet\frac{\alpha_i^{SL} + \alpha_{-i}^{SF}\psi(\kappa^\bullet)}{\alpha_i^{SL} + \hat{\alpha}_{-i}\psi(\kappa)}) = \frac{\gamma}{1 + r^\bullet\frac{\alpha_i^{SL}+\alpha_{-i}^{SF}\psi(\kappa^\bullet)}{\alpha_i^{SL}}}$$

We first show that $\gamma - \frac{1}{2}r^\bullet(\alpha_i^{SL} + \alpha_{-i}^{SF}\psi(\kappa^\bullet)) > \alpha_i^{SL}$, which we rewrite as:

$$\gamma\left(1 - \frac{1}{2}\frac{r^\bullet}{1+r^\bullet}\psi(\kappa^\bullet)\right) - \alpha_i^{SL}\left(1 + \frac{1}{2}r^\bullet\right) + \frac{1}{4}\psi(\kappa^\bullet)^2\left(\alpha_i^{SL}\frac{(r^\bullet)^2}{1+r^\bullet}\right) > 0.$$

Multiplying by $2(1+r^\bullet)$ and substituting in for α_i^{SL} gives us that this is equivalent to:

$$\frac{\gamma((2+2r^\bullet - r^\bullet\psi(\kappa^\bullet)))}{2+3r^\bullet + (r^\bullet)^2 - \frac{1}{2}\psi(\kappa^\bullet)^2(r^\bullet)^2} > \frac{\gamma(2(1+r^\bullet) - \psi(\kappa^\bullet)r^\bullet)}{2+2r^\bullet(2+r^\bullet) - \psi(\kappa^\bullet)^2(r^\bullet)^2},$$

Or:

$$r^\bullet + (r^\bullet)^2 > \frac{1}{2}\psi(\kappa^\bullet)^2(r^\bullet)^2.$$

Since $\psi(\kappa^\bullet) \leq 1$ and $r^\bullet > 0$, this inequality holds.

For the $\hat{\alpha}_{-i} \to 0$ limit, using that $\alpha_i^{SL} > 0$, to show that as $\hat{\alpha}_{-i} \to 0$, $\alpha_i^{BR}(\hat{\alpha}_{-i}; \kappa; r^\bullet \frac{\alpha_i^{SL} + \alpha_{-i}^{SF}\psi(\kappa^\bullet)}{\alpha_i^{SL} + \hat{\alpha}_{-i}\psi(\kappa)}) < \alpha_i^{SL}$ it suffices to show:

$$\gamma < \alpha_i^{SL}(1+r^\bullet) + r^\bullet\alpha_{-i}^{SF}\psi(\kappa^\bullet).$$

Substituting in for α_{-i}^{SF} shows that the right hand side is equal to:

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

Note, however, that:

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

where the right hand side is the denominator of α_i^{SL}. Therefore, it suffices to show that:

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

Multiplying both sides by $2(1+r^\bullet)$ reduces this expression to:

$$0 < \gamma\psi(\kappa^\bullet)r^\bullet,$$

which holds due to the assumptions on the parameters, thus completing the proof.

Thus, we have that there exists an equilibrium where (i) entrants assume residents play some strategy $\hat{\alpha}_{-i}$, and make inferences about $\hat{r}$ accordingly, and (ii) residents, in turn, play α_{-i}^{SF} while entrants play α_i^{SL}. We now consider the possible existence of another equilibrium, when the entrant assumes the resident's strategy is given by the previously derived $\hat{\alpha}_{-i}$. Denote:

$$\alpha_{-i}^{\bullet}(\alpha_i) := \frac{\gamma - \frac{1}{2}r^{\bullet}\psi(\kappa^{\bullet})\alpha_i}{1 + r^{\bullet}}$$

as the (rational) resident's best reply to the entrant. We have the entrant inference, in general, is:

$$\hat{r} = r^{\bullet}\frac{\alpha_i + \alpha_{-i}^{\bullet}(\alpha_i)\psi(\kappa^{\bullet})}{\alpha_i + \hat{\alpha}_{-i}\psi(\kappa)}.$$

The best reply condition yields:

$$\alpha_i + \alpha_i\hat{r} = \gamma - \frac{1}{2}\hat{r}\psi(\kappa)\hat{\alpha}_{-i}.$$

We substitute in for $\hat{r}$ and then multiply by the denominator, yielding:

$$\alpha_i(\alpha_i+\hat{\alpha}_{-i}\psi(\kappa))+\alpha_i r^{\bullet}(\alpha_i+\alpha_{-i}^{\bullet}(\alpha_i)\psi(\kappa^{\bullet})) = \gamma(\alpha_i+\hat{\alpha}_{-i}\psi(\kappa))-\frac{1}{2}r^{\bullet}(\alpha_i+\alpha_{-i}^{\bullet}(\alpha_i)\psi(\kappa^{\bullet}))\psi(\kappa)\hat{\alpha}_{-i}.$$

We consider the function:

$$\tilde{H}(\alpha_i) = \alpha_i(\alpha_i+\hat{\alpha}_{-i}\psi(\kappa))+\alpha_i r^{\bullet}(\alpha_i+\alpha_{-i}^{\bullet}(\alpha_i)\psi(\kappa^{\bullet}))-\gamma(\alpha_i+\hat{\alpha}_{-i}\psi(\kappa))+\frac{1}{2}r^{\bullet}(\alpha_i+\alpha_{-i}^{\bullet}(\alpha_i)\psi(\kappa^{\bullet}))\psi(\kappa)\hat{\alpha}_{-i}.$$

Note that this function is quadratic, which follows from inspection and the observation that $\alpha_{-i}^{\bullet}(\alpha_i)$ is linear in α_i. We also note that this function is convex, since

$$\tilde{H}''(\alpha_i) = 2(1 + r^{\bullet} - \frac{1}{2}\frac{(r^{\bullet})^2}{1 + r^{\bullet}}\psi(\kappa^{\bullet})^2) > 0,$$

where the inequality holds since $2(1 + r^{\bullet})^2 > (r^{\bullet})^2\psi(\kappa^{\bullet})^2$. We claim that it has a unique positive root. To show this, we evaluate this expression at $\alpha_i = 0$, which after substituting in for $\alpha_{-i}^{\bullet}(0) = \frac{\gamma}{1+r^{\bullet}}$, reduces to:

$$\tilde{H}(0) = \gamma(-1 + \frac{1}{2}\frac{r^\bullet}{1+r^\bullet}\psi(\kappa^\bullet))\psi(\kappa)\hat{\alpha}_{-i} < 0.$$

Therefore, since $\tilde{H}(\alpha_i)$ is a convex quadratic function which is negative at $\alpha_i = 0$, there can only be at most one positive root. It follows that given the entrant's misspecification, there is a unique equilibrium.

We now consider the setting where the entrant is dogmatic about κ and believes $r = r^\bullet$. In this case, a fixed perception about the opponent's strategy as $\hat{\alpha}_{-i}$ yields a unique best reply by Lemma 2—in particular, the entrant has a unique best reply by assumption in this case, and hence so does the resident. Furthermore, since there is no inference, it is immediate to calculate $\lim_{\hat{\alpha}_{-i}\to 0} \alpha_i^{BR}(\hat{\alpha}_{-i};\kappa;r^\bullet) = \frac{\gamma}{1+r^\bullet}$ and $\lim_{\hat{\alpha}_{-i}\to\infty} \alpha_i^{BR}(\hat{\alpha}_{-i};\kappa;r) < 0$. Since $\alpha_i^{BR}(\alpha_{-i};\kappa,r)$ is linear in α_{-i}, we therefore have a unique value of $\hat{\alpha}_{-i} > 0$ such that $\alpha_i^{BR}(\hat{\alpha}_{-i};\kappa,r^\bullet) = \alpha_i^{SL}$ since $\alpha_i^{SL} < \frac{\gamma}{1+r^\bullet}$; while this claim follows from algebra, it also follows more succinctly from the observation that $\alpha_{-i}^{SF} > 0$ and that $\alpha_i^{BR}(\alpha_{-i};\kappa,r^\bullet)$ is decreasing in α_{-i}. Thus, we have that a resident with dogmatic κ can invade against the correctly specified resident in the absence of the learning channel as well. □

A.8 Proof of Proposition 5

Proof. For the second claim, we know by Proposition 1 that for every $r^\bullet > 0$, we can find some $\bar{\kappa} > \kappa^\bullet$ so that there is a unique linear equilibrium for $\kappa_E \in (\kappa^\bullet, \bar{\kappa}]$, and furthermore entrants have strictly higher fitness than residents in this equilibrium. Take the minimum across such $\bar{\kappa}$ for the M different values of $r^m > 0$ in the M markets, and an entrant model with this value of correlation misperception strictly outperforms the residents in every market.

For the first claim, first note that in a market with true price elasticity $r^\bullet$, by the proof of Proposition 1 the correctly specified residents have correct beliefs in every linear equilibrium, and their equilibrium fitness is also uniquely determined. Also, using the subjective best response function from Lemma 2, the strategies profile (α_E, α_R) played between entrants with the perception $(\hat{r}, \hat{\kappa})$ and the residents in any linear equilibrium must satisfy $\alpha_E = \frac{\gamma - \frac{1}{2}\hat{r}\psi(\hat{\kappa})\alpha_R}{1+\hat{r}}$ and $\alpha_R = \frac{\gamma - \frac{1}{2}r^\bullet\psi(\kappa^\bullet)\alpha_E}{1+r^\bullet}$. Making the substitution for α_R in the expression for α_E, we find that α_E is uniquely determined by the linear equation

$$\alpha_E = \frac{\gamma - \frac{1}{2}\hat{r}\psi(\hat{\kappa})[\frac{\gamma - \frac{1}{2}r^\bullet\psi(\kappa^\bullet)\alpha_E}{1+r^\bullet}]}{1+\hat{r}}.$$

This shows the linear equilibrium payoff of the entrants is uniquely determined by $(r^{\bullet}, \hat{r}, \hat{\kappa})$ and is continuous in these parameters.

The goal is to find an environment with two markets with $r^1 \approx 0$, $r^2 = 2$, $\phi^1 \approx 1$, $\phi^2 \approx 0$ such that the rational resident model is resistant to invasion from any dogmatic model. Toward this end, first note that a dogmatic entrant model with the perception $\hat{r} = 0$ (and any perception of κ) always uses the strategy $\alpha_i = \gamma$ in linear equilibrium, so the resident will choose $\frac{\gamma - \frac{1}{2} 2 \cdot \psi(\kappa^{\bullet})\gamma}{3} = \gamma(\frac{1-\psi(\kappa^{\bullet})}{3})$. In a market with $r^{\bullet} = 2$, such a model gets the payoff

$$\mathbb{E}[s_i^2]\left(\frac{\gamma^2}{2} - \gamma^2\left(1 + \frac{1-\psi(\kappa^{\bullet})}{3}\right)\right).$$

This payoff is strictly negative since $\psi(\kappa^{\bullet}) > 0$. Since linear equilibrium payoffs are continuous in the entrant's perceptions $(\hat{r}, \hat{\kappa})$, we may find sufficiently small $\underline{r} > 0$ so that for every dogmatic entrant model with $\hat{r} \in [0, \underline{r}]$ and every $\hat{\kappa} \in [0, 1]$, the entrant has a strictly negative equilibrium payoff in the market with $r^{\bullet} = 2$.

Now, consider a market with $r^{\bullet} = 0$. The rational resident always chooses $\alpha_R = \gamma$ and this strategy is strictly dominant. A dogmatic entrant model with perceptions $(\hat{r}, \hat{\kappa})$ chooses $\alpha_E = \frac{\gamma - \frac{1}{2}\hat{r}\psi(\hat{\kappa})\gamma}{1+\hat{r}}$ in linear equilibrium. Find a small enough $x > 0$ so that $x\gamma - \frac{1}{2}x^2 < \frac{1}{4}\gamma^2$. Set $\bar{r} > \underline{r}$ so that $\frac{\gamma}{1+\bar{r}} = x$. For any perception $\hat{r} \geq \bar{r}$, the entrant's strategy is $\frac{\gamma - \frac{1}{2}\hat{r}\psi(\hat{\kappa})\gamma}{1+\hat{r}} \leq \frac{\gamma}{1+\hat{r}} \leq \frac{\gamma}{1+\bar{r}} = x$, so their payoff is no larger than $x\gamma - \frac{1}{2}x^2 < \frac{1}{4}\gamma^4$. This is less than half of the payoff of the rational residents, who choose strategy γ and get $\frac{1}{2}\gamma^2$.

Let $c_0 > 0$ be the rational residents' payoff when $r^{\bullet} = 0$, let $c_2 > 0$ be the rational residents' payoff when $r^{\bullet} = 2$, and let $c_s > 0$ be the Stackelberg payoff against the rational model when $r^{\bullet} = 2$. For every $r \in [\underline{r}, \bar{r}]$ and $\kappa \in [0, 1]$, let $\xi(r, \kappa)$ be the linear equilibrium payoff of the dogmatic entrant model with perceptions r, κ. We have $c_s > c_2$ but $\xi(r, \kappa) < c_0$ for every $r \in [r_0, r_1], \kappa \in [0, 1]$. So, there exists some $\epsilon_{r,\kappa} > 0$ so that $\epsilon_{r,\kappa} \cdot c_s + (1 - \epsilon_{r,\kappa}) \cdot \xi(r, \kappa) = \epsilon_{r,\kappa} \cdot c_2 + (1 - \epsilon_{r,\kappa}) \cdot c_0$ for every $r \in [r_0, r_1], \kappa \in [0, 1]$. Finally, there is some $\epsilon' > 0$ so that $\epsilon' \cdot c_s + (1 - \epsilon') \cdot (c_0/2) < \epsilon' \cdot c_2 + (1 - \epsilon') \cdot c_0$. We have that $\min_{r \in [r_0, r_1], \kappa \in [0,1]} \epsilon_{r,\kappa} > 0$ since $\xi(r, \kappa)$ is continuous, so we can find some positive $\varepsilon < \min\{\min_{r \in [r_0, r_1], \kappa \in [0,1]} \epsilon_{r,\kappa}, \epsilon'\}$ with the property that in a heterogeneous markets environment with $r^1 = 0, r^2 = 2, \phi^1 = 1 - \varepsilon, \phi^2 = \varepsilon$, the rational resident model's weighted average payoff is strictly larger than that of any dogmatic entrant model with any perception (r, κ), and in particular any dogmatic entrant model with perception $r > \bar{r}$ has weighted average payoff no larger than $\frac{3}{4}c_0$.

Finally, we can find a small enough $\delta > 0$ so that in a heterogeneous markets environment with $r^1 = \delta, r^2 = 2, \phi^1 = 1 - \varepsilon, \phi^2 = \varepsilon$, we have (i) the rational resident model's weighted average payoff is strictly larger than that of any dogmatic entrant model with any perception (r, κ) with $r \in [0, \bar{r}]$ and (ii) any dogmatic entrant model with perception $r > \bar{r}$ has weighted average payoff no larger than $\frac{3.1}{4}c_0$, which is in particular lower than the fitness of the rational resident model. $\square$

A.9 Proof of Proposition 6

For any given values of the parameters $\theta, r, \sigma_\zeta^2$, an agent's belief about the joint distribution of (s_i, ω) does not depend on κ. So, it is an equilibrium for an agent whose model satisfies the hypotheses of the proposition to use the objectively optimal strategy in the monopoly market and infer $\tilde{r}_i = r^\bullet$, $\tilde{\sigma}_{\zeta,i}^2 = (\sigma_\zeta^\bullet)^2$. We now show there is no equilibrium where the agent makes wrong inferences about r, σ_ζ^2 or chooses a different strategy. By the same argument as in the proof of Lemma 2, given the correct belief $\theta = 0$ and any beliefs about $\kappa, r, \sigma_\zeta^2$, the subjectively optimal q_i^* following the signal realization s_i is $\frac{\gamma}{1+2r}s_i$. This means in any equilibrium where the agent infers r_i, they must choose the linear strategy $\alpha_i = \frac{\gamma}{1+2r}$. Given this strategy (which in particular has $\alpha_i > 0$), the agent can set KL divergence to zero using the correct inferences $\tilde{r}_i = r^\bullet$, $\tilde{\sigma}_{\zeta,i}^2 = (\sigma_\zeta^\bullet)^2$, but KL divergence is strictly positive for any other inference. This implies beliefs are correct and strategy is optimal in any equilibrium.

A.10 Proof of Proposition 7

Consider an agent i with beliefs θ and r who has the signal realization s_i. Their expected payoff from choosing quantity q_i is:

$$q_i \mathbb{E}[\theta + \omega - rq_i + \zeta \mid s_i] - \frac{1}{2}q_i^2 = q_i \cdot [\theta + \gamma s_i - rq_i] - \frac{1}{2}q_i^2$$

Taking FOC in q_i, we find that the subjectively optimal q_i^* following the signal realization s_i is $\frac{\theta}{1+2r} + \frac{\gamma}{1+2r}s_i$. (The second-order condition is satisfied provided $r > -1/2$.) If the agent has dogmatic misperception of $r \geq 0$, then their equilibrium strategy has a wrong slope and they do not play the optimal q_i for any (except possibly one) s_i. If the agent has a misperception of θ, then for any inference $r \geq 0$, their equilibrium strategy either has a wrong intercept or a wrong slope, so they again do not play the optimal q_i for any (except possibly one) s_i. This shows that in any equilibrium, a model with a fixed and wrong r or a fixed and wrong θ must lead to a strict loss in expected utility.

A.11 Proof of Proposition 8

Proof. Consider the society where $\Theta_R = \Theta_E = \Theta(\kappa^\bullet)$, where Θ_R is the resident model and Θ_E is the entrant model. For any linear equilibrium with behavior $(\sigma_{R\to R}, \sigma_{R\to E}, \sigma_{E\to R}, \sigma_{E\to E})$ and beliefs $\mu_R \in \Delta(\Theta_R)$ and $\mu_E \in \Delta(\Theta_E)$, there exists another linear equilibrium $(\sigma'_{R\to R}, \sigma'_{R\to E}, \sigma'_{E\to R}, \sigma'_{E\to E})$ where $\sigma'_{g,g'} = \sigma_{R\to R}$ for all $g, g' \in \{A, B\}$ and all agents have the belief μ_R. The uniqueness of linear equilibrium from Assumption 1 implies $\alpha_{R\to E}(\kappa^\bullet) = \alpha_{E\to R}(\kappa^\bullet) = \alpha_{E\to E}(\kappa^\bullet) = \alpha_{R\to R}(\kappa^\bullet) = \alpha^\bullet$.

Now consider the society where $\Theta_E = \Theta(\kappa)$. By assumption, there exists a linear equilibrium where $\alpha_{R\to R}(\kappa) = \alpha_{R\to R}(\kappa^\bullet)$. Since we also take it to be unique, we must in fact have $\alpha_{R\to R}(\kappa) = \alpha_{R\to R}(\kappa^\bullet)$ for all κ, so the fitness of model $\Theta(\kappa^\bullet)$ in the unique linear equilibrium is

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

Given λ and entrant model $\Theta(\kappa)$, the entrant's fitness in the unique linear equilibrium is

$$\mathbb{E}^\bullet\left[\mathbb{E}^\bullet\left[(1-\lambda)u_1^\bullet(\alpha_{E\to R}(\kappa)s_1, \alpha_{R\to E}(\kappa)s_2, \omega) + (\lambda)u_1^\bullet(\alpha_{E\to E}(\kappa)s_1, \alpha_{E\to E}(\kappa)s_2, \omega) \mid s_1\right]\right].$$

Differentiate and evaluate at $\kappa = \kappa^\bullet$. At $\kappa = \kappa^\bullet$, agents with models Θ_R and Θ_E have the same fitness since they play the same strategies. So, a non-zero sign on the derivative would give the desired susceptibility to invasion from models with either slightly higher or slightly lower κ. This derivative is:

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

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. □

B More General Population Sizes and Matching Assortativities

In this section, we discuss an extension to the baseline environment where entrants can make up a positive share of the population and matching can be neither perfectly uniform nor perfectly assortative. One complication with this more general environment is that no linear equilibrium as we defined earlier in this paper may exist, since there may not be any belief about the free parameters that fully explains the distribution of consequences observed in matches against both groups of opponents. Instead, we adopt the solution concept of *equilibrium zeitgeist* from He and Libgober (2024), where agents make the best-fitting inference in terms of minimizing Kullback–Leibler divergence (KL divergence).

We show that in the special case where the population share of the entrants is zero and the probability of matching uniformly is either one or zero, the equilibrium zeitgeists in this more general environment exactly correspond to the linear equilibria in the baseline environment with uniform or assortative matching, respectively. We also show that provided the share of entrants is sufficiently near zero and matching assortativity is sufficiently close to uniform matching or assortative matching, the same results about susceptibility and resistance to invasion from the baseline model continue to hold. That is, $\kappa^\bullet$ is susceptible to invasion from projection bias and resistant to invasion from correlation neglect, when the share of entrants is sufficiently small and matching is sufficiently near uniform matching. On the other hand, $\kappa^\bullet$ is susceptible to invasion from correlation neglect and resistant to invasion from projection bias, when the share of entrants is sufficiently small and matching is sufficiently near perfectly assortative matching.

We consider the same stage game and information structure as in Section 2.1. We restriction attention to linear strategies and let the space of strategies by parametrized by $\mathbb{A} = [0, \bar{M}_\alpha]$ for $\bar{M}_\alpha < \infty$, where $\alpha_i \in [0, \bar{M}_\alpha]$ refers to the strategy that chooses quantity $\alpha_i s_i$ after every signal s_i. We view each model as a set Θ of feasible parameter values. For an agent with a dogmatic model, the model is a singleton set $\Theta = \{(\tilde{\kappa}, \tilde{r}, \tilde{\sigma}_\zeta^2)\}$. For an agent with a flexible model who thinks signal correlation is equal to $\tilde{\kappa}$, the model is given by $\Theta = \{(\tilde{\kappa}, r, \tilde{\sigma}_\zeta^2) : r \in [0, \bar{M}_r], \sigma_\zeta^2 \in [0, \bar{M}_{\sigma_\zeta^2}]\}$ for some $\bar{M}_r, \bar{M}_{\sigma_\zeta^2} < \infty$.

We have assumed that the space of feasible linear strategies $\alpha_i \in [0, \bar{M}_\alpha]$ and the domain of inference are compact to guarantee the existence of our solution concept, to be introduced below. For some of our results, we use 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 with $(\sigma_\zeta^\bullet)^2 \geq L_1$ and with models where $r \in [0, \bar{M}_r]$, $\sigma_\zeta \in [0, \bar{M}_{\sigma_\zeta}]$ 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.

We consider a society with two groups, entrants (E) and residents (R). There is $p_E \in [0,1]$ share of entrants and $p_R = 1 - p_E$ share of residents in the population. Entrants have model Θ_E and incumbents have model Θ_R. We also fix a matching assortativity parameter $\lambda \in [0,1]$. Each agent of group g is matched with someone in their group with probability $\lambda + (1-\lambda)p_g$, and matched with someone from the other group with the complementary probability. The setup in the main text of the paper corresponds to the special case of $p_E = 0$, with $\lambda = 0$ (uniform matching) or $\lambda = 1$ (assortative matching).

We use the solution concept of equilibrium zeitgeist from He and Libgober (2024). We provide its definition below, specialized to the setting of this particular game.

Definition A.1. For fixed population share (p_E, p_R) and assortativity λ, an *equilibrium zeitgeist* for this game consists of strategies $a_{EE}, a_{ER}, a_{RE}, a_{RR}$ and beliefs $\mu_E \in \Theta_E$ and $\mu_R \in \Theta_R$, such that:

- For each group g and opponent group g', the strategy $a_{g,g'}$ is a subjective best response against $a_{g',g}$ given the belief μ_g.
- For each group g, the belief μ_g solves

$$\underset{\mu \in \Theta_g}{\arg\min} \left\{ (\lambda + (1-\lambda)p_g) \cdot K(\mu; a_{g,g}, a_{g,g}) + (1-\lambda)(1-p_g) \cdot K(\mu; a_{g,-g}, a_{-g,g}) \right\}$$

  where $-g$ is the group other than g and $K(\mu; a, a')$ refers to the KL divergence from the expected distribution of consequences under parameters μ to the objective distribution of consequences, under the strategy profile (a, a').

Next, we show that equilibrium zeitgeist is an appropriate extension of linear equilibria into a setting with a non-zero population share of entrants and more general assortativity by showing that these two solution concepts coincide when $p_E = 0$, $\lambda = 0$ or $\lambda = 1$. Technically, a linear equilibrium does not require us to specify strategies used in all types of matches (we only specify three strategies for linear equilibrium with uniform matching, and only two strategies for linear equilibrium with assortative matching). To say that the set of linear equilibrium is equivalent to the set of equilibrium zeitgeists, we formally mean that every equilibrium zeitgeist is a linear equilibrium with the appropriate strategies removed, and that every linear equilibrium can be augmented with some strategies to become an equilibrium zeitgeist.

We assumed the compactness of the strategy space and the inferences space to ensure the existence of equilibrium zeitgeists, but we will need to make these compact spaces large enough to avoid running into corner solutions.

Lemma A.1. *Suppose $p_A = 0$ and either $\lambda = 0$ or $\lambda = 1$. With high enough price volatility and large enough strategy space and inference space, the set of equilibrium zeitgeists is equivalent to the set of linear equilibria with uniform matching or assortative matching, respectively.*

Proof. Let $\bar{M}_\alpha \geq \frac{1/\sigma_\epsilon^2}{1/\sigma_\epsilon^2 + 1/\sigma_\omega^2}$ be given. Let L_1 be defined as L in the proof of Lemma 3. Let L_2 be $\max_{\kappa \in [0,1], 0 \leq \alpha_i, \alpha_{-i} \leq \gamma} \mathrm{Var}_\kappa[\omega - \frac{1}{2} r^\bullet \alpha_{-i} s_{-i} \mid s_i] + (\sigma_\zeta^\bullet)^2$, which, by the arguments in the proof of Lemma 3, is finite and independent of s_i. And let L_3 be $\max_{\kappa \in [0,1], 0 \leq \alpha_i, \alpha_{-i} \leq \gamma} r^\bullet \frac{\alpha_i + \alpha_{-i}\psi(\kappa^\bullet)}{\alpha_i + \alpha_{-i}\psi(\kappa)}$. We show that with these choices of L_1, L_2, L_3, the set of equilibrium zeitgeists is equivalent to the set of linear equilibria.

In any equilibrium zeitgeist, the subjective best response (even if there were no restrictions on the strategy) to any opponent's strategy must be a linear strategy with slope in $[0, \gamma]$, by Lemma 2. Since $\bar{M}_\alpha \geq \gamma$, such best responses are feasible strategies. So, in every equilibrium zeitgeist, we must have $\alpha_i, \alpha_{-i} \in [0, \gamma]$. By the choice of L_2 and L_3 and using Lemma 3, there is a belief that is self-confirming within the model. Hence, all equilibrium zeitgeists must have self-confirming beliefs. Thus, all equilibrium zeitgeists are linear equilibria.

Conversely, any linear equilibrium must have $\alpha_i, \alpha_{-i} \in [0, \gamma]$, so they satisfy $\alpha_i, \alpha_{-i} \leq \bar{M}_\alpha$ in the bounded strategy space. The inferences are within the bounds $[0, \bar{M}_r]$ and $[0, \bar{M}_{\sigma_\zeta}]$ and are self-confirming, so their KL divergence cannot be further reduced. So linear equilibria are equilibrium zeitgeists. □

Finally, we show that Proposition 1 and Proposition 2 are robust to small perturbations in population size and matching assortativity. A useful lemma is that the set of equilibrium zeitgeists (for any values of $\bar{M}_\alpha, \bar{M}_{\sigma_\zeta}, \bar{M}_r < \infty$) is non-empty and upper hemicontinuous in p_E and λ.

Lemma A.2. *For any fixed $\bar{M}_\alpha, \bar{M}_{\sigma_\zeta}, \bar{M}_r < \infty$, the set of equilibrium zeitgeists in this game is non-empty and upper hemicontinuous with respect to p_E and λ.*

Proof. Propositions A.1 and A.2 from He and Libgober (2024) provide a set of conditions on the stage game and the models that guarantee the existence and upper hemicontinuity of the set of equilibrium zeitgeists.

Assumption A.1. *$\mathbb{A}$, Θ_E, and Θ_R are compact, metrizable spaces.*

For each group g and $\mu \in \Theta_g$, let $U_g(a, a'; \mu)$ be subjective expected utility under belief μ and strategy profile (a, a').

Assumption A.2. *U_E, U_R are continuous.*

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

Under Assumption A.3, we have the well-defined functions $K_E : \Theta_E \times \mathbb{A}^2 \to \mathbb{R}_+$ and $K_R : \Theta_R \times \mathbb{A}^2 \to \mathbb{R}_+$, where $K_g(\mu; a_i, a_{-i}) = K(\mu; a_i, a_{-i})$.

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 $\mu \in \Theta_E \cup \Theta_R$, $a_g \mapsto U_g(a_g, a_{-g}; \mu)$ is quasiconcave.*

We now show that these assumptions hold.

Assumption A.1 holds as $\mathbb{A}, \Theta_E, \Theta_R$ 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 $(\hat{r}, \kappa, \sigma_\zeta^2)$ is $\mathbb{E}[s_i^2] \cdot \left(\alpha_i \gamma - \frac{1}{2}\hat{r}\alpha_i^2 - \frac{1}{2}\hat{r}\psi(\kappa)\alpha_i\alpha_{-i} - \frac{1}{2}\alpha_i^2\right)$. 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 $\zeta^\bullet$ 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 $\alpha_{-i}, \hat{r}$ 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 a upper bound for choices of the parameters $\alpha_{-i}, \hat{r}$ 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

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

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_E, K_R 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. □

Proposition A.1. *Fix $r^\bullet > 0$, $\kappa^\bullet \in [0, 1]$. For any given $\kappa_E \neq \kappa_R$, with high enough price volatility and large enough strategy space and inference space, there exists $\epsilon > 0$ with the property that for any $p_E \leq \epsilon$ and any $\lambda \geq 1 - \epsilon$, in every equilibrium zeitgeist with flexible models Θ_E and Θ_R the residents have strictly higher expected utility than the entrants if $\kappa_E < \kappa_R$ and the residents have strictly lower expected utility than the entrants if $\kappa_E > \kappa_R$.*

Also, with high enough price volatility and large enough strategy space and inference space, *there exist $\underline{\kappa} < \kappa^\bullet < \bar{\kappa}$ so that:*

1. *For any $\kappa \in [\underline{\kappa}, \kappa^\bullet)$, we may find $\epsilon > 0$ with the property that for any $p_E \leq \epsilon$ and any $\lambda \leq \epsilon$, in every equilibrium zeitgeist with flexible models Θ_E and Θ_R where $(\kappa_R, \kappa_E) = (\kappa^\bullet, \kappa)$, the residents have strictly higher expected utility than the entrants.*

2. *For any $\kappa \in (\kappa^\bullet, \bar{\kappa}]$, we may find $\epsilon > 0$ with the property that for any $p_E \leq \epsilon$ and any $\lambda \leq \epsilon$, in every equilibrium zeitgeist with flexible models Θ_E and Θ_R where $(\kappa_R, \kappa_E) = (\kappa^\bullet, \kappa)$, the residents have strictly lower expected utility than the entrants.*

Proof. For the first part, set $p_E = 0$ and $\lambda = 1$. For any $\bar{M}_\alpha \geq \frac{1/\sigma_\epsilon^2}{1/\sigma_\epsilon^2 + 1/\sigma_\omega^2}$, by Lemma A.1 find L_1, L_2, L_3. For any $(\sigma_\zeta^\bullet)^2 \geq L_1$, $\bar{M}^2_{(\sigma_\zeta)} \geq (\sigma_\zeta^\bullet)^2 + L_2$ and $\bar{M}_r \geq L_3$, Lemma A.1 implies that the set of equilibrium zeitgeists is equivalent to the set of linear equilibria. But the proof of Proposition 2 shows that in every such linear equilibrium, the residents have strictly higher expected utility than the entrants if $\kappa_E < \kappa_R$ and the residents have strictly lower expected utility than the entrants if $\kappa_E > \kappa_R$. The same strict inequality then holds for the set of equilibrium zeitgeists. Since the set of equilibrium zeitgeists is upper hemicontinuous in p_E and λ by Lemma A.2, the same must also be true for p_E close enough to 0 and λ close enough to 1.

For the second part, set $p_E = 0$ and $\lambda = 0$. For any $\bar{M}_\alpha \geq \frac{1/\sigma_\epsilon^2}{1/\sigma_\epsilon^2 + 1/\sigma_\omega^2}$, by Lemma A.1 find L_1, L_2, L_3. For any $(\sigma_\zeta^\bullet)^2 \geq L_1$, $\bar{M}^2_{(\sigma_\zeta)} \geq (\sigma_\zeta^\bullet)^2 + L_2$ and $\bar{M}_r \geq L_3$, Lemma A.1 implies that the set of equilibrium zeitgeists is equivalent to the set of linear equilibria. Use the values of $\underline{\kappa}, \bar{\kappa}$ from Proposition1. The proof of Proposition 1 shows that for any $\kappa \in [\underline{\kappa}, \kappa^\bullet)$, in the unique linear equilibrium with $(\kappa_R, \kappa_E) = (\kappa^\bullet, \kappa)$ the residents have strictly higher expected utility than the entrants. The same strict inequality then holds for the set of equilibrium zeitgeists. Since the set of equilibrium zeitgeists is upper hemicontinuous in p_E and λ by Lemma A.2, the same must also be true for p_E close enough to 0 and λ close enough to 0. The analogous argument applies for the case of $\kappa \in (\kappa^\bullet, \bar{\kappa}]$. □

References

Alger, I. and J. Weibull (2013): "Homo Moralis-Preference Evolution Under Incomplete Information and Assortative Matching," *Econometrica*, 81, 2269–2302.

——— (2019): "Evolutionary models of preference formation," *Annual Review of Economics*, 11, 329–354.

Angeletos, G.-M. and A. Pavan (2007): "Efficient use of information and social value of information," *Econometrica*, 75, 1103–1142.

Asker, J., C. Fershtman, and A. Pakes (2023): "The Impact of AI Design on Pricing," *Journal of Economics and Management Strategy*, 1–29.

Bergemann, D., T. Heumann, and S. Morris (2017): "Information and Interaction," Working paper.

Bergemann, D. and S. Morris (2013): "Robust predictions in games with incomplete information," *Econometrica*, 81, 1251–1308.

Berman, R. and Y. Heller (2024): "Naive Analytics: The Strategic Advantage of Algorithmic Heuristics," *Working Paper*.

Calvano, E., G. Calzolari, V. Denicoló, and S. Pastorello (2020): "Artificial Intelligence, Algorithmic Pricing, and Collusion," *American Economic Review*, 110, 3267–3297.

Dekel, E., J. Ely, and O. Yilankaya (2007): "Evolution of preferences," *Review of Economic Studies*, 74, 685–704.

Esponda, I. and D. Pouzo (2016): "Berk–Nash equilibrium: A framework for modeling agents with misspecified models," *Econometrica*, 84, 1093–1130.

Fershtman, C. and K. L. Judd (1987): "Equilibrium Incentives in Oligopoly," *American Economic Review*, 77, 927–940.

Frick, M., R. Iijima, and Y. Ishii (2024): "Welfare comparisons for biased learning," *American Economic Review*, 1612–1649.

Friedman, M. (1953): *Essays in Positive Economics*, University of Chicago Press.

Fudenberg, D. and G. Lanzani (2023): "Which misspecifications persist?" *Theoretical Economics*, 18, 1271–1315.

Gagnon-Bartsch, T., M. Pagnozzi, and A. Rosato (2021): "Projection of Private Values in Auctions," *American Economic Review*, 111, 3256–3298.

Gagnon-Bartsch, T. and A. Rosato (2024): "Quality is in the eye of the beholder: taste projection in markets with observational learning," *American Economic Review*, 114, 3746–3787.

Hansen, K., K. Misra, and M. Pai (2021): "Frontiers: Algorithmic collusion: Supra-competitive prices via independent algorithms," *Marketing Science*, 40, 1–12.

He, K. and J. Libgober (2024): "Misspecified Learning and Evolutionary Stability," *Working Paper.*

Heidhues, P., B. Koszegi, and P. Strack (2018): "Unrealistic expectations and misguided learning," *Econometrica*, 86, 1159–1214.

Heifetz, A., C. Shannon, and Y. Spiegel (2007): "What to maximize if you must," *Journal of Economic Theory*, 133, 31–57.

Madarász, K. (2012): "Information Projection: Model and Applications," *Review of Economic Studies*, 79, 961–985.

Miyashita, M. and T. Ui (2023): "LQG Information Design," *Working Paper.*

Robson, A. J. and L. Samuelson (2011): "The evolutionary foundations of preferences," in *Handbook of Social Economics*, Elsevier, vol. 1, 221–310.

Vives, X. (1988): "Aggregation of information in large Cournot markets," *Econometrica*, 851–876.