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

Kevin He[†] Jonathan Libgober[‡]

August 26, 2025

Abstract

We study misperceptions of private-signal correlation in an incomplete-information Cournot duopoly game. Exaggerating the correlation between players' demand signals is beneficial when agents hold flexible beliefs about price elasticity, but harmful when their beliefs are dogmatically correct. For agents with flexible beliefs who 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 formalize this idea in the context of a linear-quadratic-normal (LQN) Cournot duopoly game of incomplete information, similar to Vives (1988). 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 (2025); 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, with residents making up almost the entire population and entrants comprising a negligible fraction. Price elasticity is a structural parameter of the market that remains fixed across games, but idiosyncratic demand shocks cause the demand intercept to be drawn i.i.d. across games. Residents have a fixed perception about how demand signals of the two game participants are correlated, while entrants may have a different perception. Our leading performance criterion posits that the entrants' model is favored if their payoffs in matches against the residents are higher than the residents' payoffs when they play against each other.[1] We show that if agents correctly know the persistent price elasticity, then this criterion favors entrants who underestimate correlation in different players' signals. However, this conclusion is reversed if agents do not know price elasticity and infer this parameter from game outcomes. In other words, a given error in higher-order beliefs about the transient state may benefit agents if and only if they 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 performance implications of projection bias 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 (at least to some extent) 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 direct effect of projection bias thus harms the agent's profits.

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

---

[1] In the literature on the indirect evolutionary approach discussed in Section 1.2, this setup is known as *uniform matching.* An alternative, *assortative matching*, considers entrants and residents playing among their respective groups and not interacting with each other. We briefly discuss this contrast in Section 5.3.

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 indirect effect of projection bias is the opposite of its direct effect.

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 correctly know the price elasticity ex-ante with certainty.

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. The notion that subjective best replies may be endogenously determined when agents interpret outcomes through the lens of a misspecified model is in the spirit of recent contributions on misspecified Bayesian learning (Esponda and Pouzo (2016); Frick, Iijima, and Ishii (2024); Heidhues, Koszegi, and

Strack (2018), among others). These exercises typically study the implications of mislearning persistent parameters on behavior and welfare, taking the misspecifications as exogenously given.

Our companion paper, He and Libgober (2025), develops a general framework applying the indirect evolutionary approach to accommodate the selection of misspecified models. As evolutionary frameworks have a long tradition of providing formal payoff-based foundations for departures from rational utility maximization, applying the same type of framework to the context of model selection provides a way to endogenize misspecifications using a familiar channel. While He and Libgober (2025) justify the inferences underlying the equilibrium concept used here, it leaves open the question of which *specific* biases might be viable. By contrast, the present paper studies a particular misspecification, motivated by the difficulties of forming correct higher-order beliefs, in a tractable linear-quadratic-normal (LQN) setting. Crucially, this paper's main message—that the welfare consequences of a given higher-order misperception may depend on agents' knowledge about the persistent game parameters—relies on the structure of the environment. Models in He and Libgober (2025)'s general setup cannot always be "factored" into fixed parameters and free parameters that represent beliefs about multiple orthogonal dimensions, like the particular models considered in this paper's environment. Given the different focuses, the papers are also quite distinct methodologically, with the present paper's analysis making heavy use of explicit formulas for Gaussian inference and other properties of LQN games.

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 hold even when this assumption is relaxed). Our framework similarly assumes that agents' perceptions of signal correlation are observable by others, which provides scope for departures from rationality (as in these other works). Our results also rely on agents correctly knowing others' strategies in equilibrium; we discuss the implications of relaxing this assumption in Section 5.1.

Much 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 once players can condition their actions on their private signals. Other papers 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 algorithms are misspecified (as they often must be given the complexity of the market environment), then they can converge to excessively collusive prices. For example, Calvano, Calzolari, Denicoló, and Pastorello (2020); Hansen, Misra, and Pai (2021); Asker, Fershtman, and Pakes (2023) 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 misperception 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).[2] 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] Of course, in practice these pricing algorithms may also exhibit first-order misperceptions about the market environment. Section 5.2.2 provides a justification for why such errors are more likely to be eliminated: if the seller is sometimes a monopolist and sometimes a duopolist, then a pricing algorithm with a first-order misperception cannot improve their profit in duopoly markets without lowering their profit in monopoly markets, but an algorithm with a second-order misperception can do so.

2 Framework

Following the indirect evolutionary approach and our companion paper He and Libgober (2025), 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 distribution of the error term ϵ_i is specified below.) 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 let:

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

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,[3] 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

---

[3] 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.

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 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). We do not require $-i$ to also derive beliefs from the same interaction, since part of our analysis concerns cases where $-i$'s beliefs are primarily shaped by the consequences they observe in other matches—specifically, when i belongs to a negligible entrant population.

Definition 3. A strategy profile Q_i, Q_{-i} and belief profile $(\tilde{\kappa}_i, \tilde{r}_i, \tilde{\sigma}_{\zeta,i}^2)$, $(\tilde{\kappa}_{-i}, \tilde{r}_{-i}, \tilde{\sigma}_{\zeta,-i}^2)$ 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 $(\tilde{\kappa}_k, \tilde{r}_k, \tilde{\sigma}_{\zeta,k}^2)$.
- For the first player i, $\tilde{\kappa}_i$ is the fixed parameter given by i's model, and $(\tilde{r}_i, \tilde{\sigma}_{\zeta,i}^2)$ 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*. For most of the paper, we assume that 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 Section 5.3, we discuss how our results change with assortative matching, where agents are always matched within the group that uses the same model.

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. 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 if there exists at least one linear equilibrium and $U^\bullet(\alpha_{R\to R}, \alpha_{R\to R}) \geq U^\bullet(\alpha_{E\to R}, \alpha_{R\to E})$ in every linear equilibrium, and we say κ_R is *susceptible to invasion* if there exists at least one linear equilibrium and $U^\bullet(\alpha_{R\to R}, \alpha_{R\to R}) < U^\bullet(\alpha_{E\to R}, \alpha_{R\to E})$ 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 of the opponent. In particular, in an environment where residents have flexible models, their 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.

Errors in higher-order beliefs have received significant attention in the behavioral literature recently (e.g., Gagnon-Bartsch et al. (2021); Gagnon-Bartsch and Rosato (2024); Madarász (2012)). In particular, agents may find them harder to recognize than errors in first-order beliefs, and thus they may be more likely to emerge in the first place. We formalize one version of the claim that higher-order errors are "harder to recognize" 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.[4] In our setting, we could justify such an assumption by imagining that agents

---

[4] 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).

sample a large number of others from each of the resident and entrant populations and observe 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[5] in the stage game. If entrants can flexibly make self-confirming misinferences about others' strategies, then there are always multiple equilibria given any misperception of signal correlation, with the entrants doing strictly better than the residents in some equilibria but weakly worse in others. So once again, we would not be able to select among different correlation perceptions on the basis of payoffs.

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:*

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

---

[5] 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.

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, and in particular the explicit formulas for Gaussian Bayesian updating. The coefficient γ that characterizes 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, belief about 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} perceived correlation $\kappa \in [0,1]$ and belief about elasticity $r \geq 0$, 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 their private information about market demand when their 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 and beliefs. 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 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. First, suppose agents have flexible models.

Proposition 1. *Fix any $r^\bullet > 0$, $\kappa^\bullet \in [0,1]$ and $(\sigma_\zeta^\bullet)^2 \geq L$. Assume all agents have flexible models.*

(a) If $2 + r^\bullet > \psi(\kappa^\bullet)(2 + r^\bullet + 2\psi(\kappa^\bullet)r^\bullet)$, there exist some $\kappa^* \in (\kappa^\bullet, 1)$ and a neighborhood $N(\kappa_R)$ around every $\kappa_R \in [0,1]$ such that:

- *If $\kappa_R < \kappa^*$, then κ_R is susceptible to invasion from any $\kappa_E \in N(\kappa_R)$ with $\kappa_E > \kappa_R$ and resistant to invasion from any $\kappa_E \in N(\kappa_R)$ with $\kappa_E \leq \kappa_R$.*
- *If $\kappa_R > \kappa^*$, then κ_R is susceptible to invasion from any $\kappa_E \in N(\kappa_R)$ with $\kappa_E < \kappa_R$ and resistant to invasion from any $\kappa_E \in N(\kappa_R)$ with $\kappa_E \geq \kappa_R$.*
- *If $\kappa_R = \kappa^*$, then κ_R is resistant to invasion from any $\kappa_E \in N(\kappa^*)$.*

(b) If $2 + r^\bullet \leq \psi(\kappa^\bullet)(2 + r^\bullet + 2\psi(\kappa^\bullet)r^\bullet)$, there exists a neighborhood $N(\kappa_R)$ around every $\kappa_R \in [0,1]$ such that κ_R is susceptible to invasion from any $\kappa_E \in N(\kappa_R)$ with $\kappa_E > \kappa_R$ and resistant to invasion from any $\kappa_E \in N(\kappa_R)$ with $\kappa_E \leq \kappa_R$.

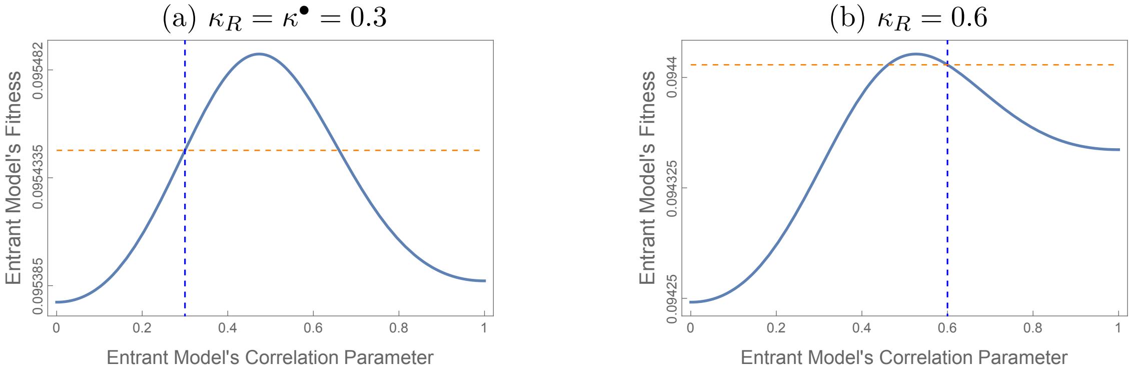

(a) $\kappa_R = \kappa^\bullet = 0.3$

(b) $\kappa_R = 0.6$

Figure 1: Fitness of flexible entrant models with different correlation perceptions when residents have flexible models 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. In the plot on the left, the residents are correctly specified. In the plot on the right, the residents are also misspecified with $\kappa_R > \kappa^*$.

In particular, when $\kappa_R = 1$, the resident model is resistant to invasion from any κ_E in the neighborhood $N(1) \subseteq [0, 1]$.

Specializing to the case where residents are correctly specified ($\kappa_R = \kappa^\bullet$), Proposition 1 tells us that these residents are susceptible to invasion by projection-biased entrants. 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 credibly respond to their opponents' play in an overly aggressive manner.[6] The rational residents back down and yield a larger share of the surplus. Since residents with the correctly specified model always correctly infer price elasticity in equilibrium, the same result about projection bias also obtains if the residents have dogmatic models with objectively correct fixed parameters (instead of flexible models with the correct correlation perception).

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

---

[6] 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.

suboptimal production. Figure 1a illustrates this non-monotonicity of fitness in κ_E. Here, we present a numerical example of the entrant's fitness as a function of their correlation perception, when facing correctly specified residents with $\kappa_R = \kappa^\bullet = 0.3$. 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, increasing the entrant's correlation perception κ_E is only locally beneficial: there is a range of 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 also fully characterizes susceptibility and resistance to such "local" mutations in correlation perception for misspecified residents with flexible models. The proof of Proposition 1 shows that the entrants' equilibrium strategy is locally increasing in κ_E when κ_E is close enough to κ_R. So, the question of resistance to invasion boils down to whether the entrants would benefit from choosing a slightly more aggressive or slightly less aggressive strategy compared to the residents.

When $2 + r^\bullet > \psi(\kappa^\bullet)(2 + r^\bullet + 2\psi(\kappa^\bullet)r^\bullet)$, there is some interior value of correlation perception $\kappa^* > \kappa^\bullet$ such that a slightly more aggressive strategy is beneficial when $\kappa_R < \kappa^*$ and a slightly less aggressive strategy is beneficial when $\kappa_R > \kappa^*$. The idea is that if the residents were to exhibit a significant amount of projection bias themselves (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. Residents with low enough κ_R are susceptible to invasion from entrants with projection bias, while those with high enough κ_R are susceptible to invasion from entrants with correlation neglect. This is illustrated in Figure 1b: when the resident has the correlation misperception $\kappa_R = 0.6$ (which is above κ^* in that environment), entrants with κ_E slightly lower than 0.6 do better than the residents. Residents with $\kappa_R = \kappa^*$ are resistant to invasion from all nearly κ_E. In fact, the proof of Proposition 1 implies such residents are also resistant to entrants with *any* kind of model (even those not in the class of models discussed here), provided at least one linear equilibrium exists. While our framework does not formally study the long-run distribution of biases in a world with a rich set of potential flexible models, the intuition here is that the model with correlation misperception $\kappa_R = \kappa^* > \kappa^\bullet$ may be the ones that eventually prevail because of their robustness to invasion.

When $2 + r^\bullet \leq \psi(\kappa^\bullet)(2 + r^\bullet + 2\psi(\kappa^\bullet)r^\bullet)$, we instead find that a slightly more aggressive strategy always improves the entrants' payoffs. So in particular, residents with $\kappa_R = 1$ are resistant to invasion from all nearby κ_E, since the entrants cannot have a correlation perception above 1.

Proposition 1 considers the selection of correlation misperception with flexible models. But as discussed before, the direct effect and the indirect effect of correlation misperception go in opposite directions. The next result shows that if entrants have dogmatic models so that the learning channel is shut down, the conclusions of Propositions 1 are exactly reversed. In Proposition 2, we imagine that the entrants initially belonged to the same community as the residents (who have flexible models), and thus shared the same correlation perception and the same belief about price elasticity (denoted r_R). Then, the entrants adopt a new *dogmatic* model with a different elasticity perception κ_E, while keeping their previous belief r_R about price elasticity as a fixed parameter. This exercise lets us isolate the direct effect of changing the correlation perception on equilibrium payoffs.

Proposition 2. *Fix any $r^\bullet > 0$, $\kappa^\bullet \in [0,1]$ and $(\sigma_\zeta^\bullet)^2 \geq L$. Suppose residents have flexible models with correlation perception $\kappa_R \in [0,1]$. Suppose entrants have dogmatic models with fixed parameters $\kappa_E \in [0,1]$ and $r_E = r_R$, where r_R is the unique equilibrium belief about price elasticity in the resident population induced by κ_R. There is a unique linear equilibrium and the equilibrium strategy $\alpha_{E\to R}$ is strictly decreasing in κ_E. So, all of Proposition 1's statements about susceptibility and resistance to invasion from local, directional mutations in entrant correlation perception κ_E hold with reversed directions.*

Consider first the case where the residents are correctly specified, so $\kappa_R = \kappa^\bullet$ and $r_R = r^\bullet$. We previously found that these residents are susceptible to invasion from entrants with flexible models who exhibit a small amount of projection bias. Proposition 2 says when entrants have dogmatic models and know the true elasticity, the residents are instead resistant to invasion from entrants with projection bias. That is, the indirect effect of elasticity misinference is responsible for the evolutionary advantage conferred by the misperceptions of κ analyzed in Propositions 1. Intuitively, Proposition 2 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. As we saw in the previous two propositions, this force is 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.

Proposition 2 also applies to misspecified residents. Recall that the proof of Proposition 1 fully characterizes when the entrants benefit from committing to a slightly more (or slightly less) aggressive strategy. This characterization continues to hold, but the equilibrium strategy for entrants with dogmatic models is a decreasing function of their correlation perception κ_E (whereas it is an increasing function of κ_E for entrants with flexible models.) This reverses the conclusions of Proposition 1. For example, if $2 + r^\bullet > \psi(\kappa^\bullet)(2 + r^\bullet + 2\psi(\kappa^\bullet)r^\bullet)$, then there exists some $\kappa^* > \kappa^\bullet$ such that residents with any correlation perception $\kappa_R < \kappa^*$ are susceptible to invasion from entrants with dogmatic models that stipulate a slightly *lower* correlation perception (assuming the fixed parameter for price elasticity in the dogmatic entrant model is equal to the equilibrium elasticity belief of the residents).

In summary, 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. (The only exception here is for residents with $\kappa_R = \kappa^*$, who are resistant to all local mutations in correlation perception regardless of whether the entrants can update their inference about price elasticity.)

5 Extensions

5.1 Incorrect Beliefs about 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 first consider the case of dogmatic strategy *misperception* before allowing for strategies to be flexibly *misinferred*.

5.1.1 Strategy Misperception

We describe how to modify our solution concept to accommodate misperceptions of others' strategies. For simplicity, we restrict attention to the case of correctly specified residents (who correctly know others' strategies in equilibrium). We suppose that entrants are endowed with both a dogmatic correlation misperception κ and a dogmatic misperception of residents' strategy $\hat{\alpha}_R$, which may be different from the residents' actual equilibrium strategy $\alpha_R^\bullet$. We consider both the case where entrants have flexible models, as well as the case where entrants have dogmatic models with fixed parameters $\tilde{r} = r^\bullet$ and $\tilde{\sigma}_\zeta^2 = (\sigma_\zeta^\bullet)^2$.

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 $\tilde{r}_E, \tilde{\sigma}_{\zeta,E}^2$ to rationalize the market price distribution under the hypothesis that residents use the strategy $\hat{\alpha}_R$.

To find the inferred elasticity, consider an entrant who misperceives signal correlation as κ and misperceives the opponent's strategy as $\hat{\alpha}_{-i}$, when in fact the opponent's strategy is $\alpha_{-i}^\bullet$. Suppose the entrant plays α_i. The same arguments as in the proof of Lemma 3 then imply that the entrant will infer elasticity to be:

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

This expression suggests that misperceptions of κ and misperceptions of others' strategies may have similar effects on inference and behavior. The next result shows that the *combination* of these two misperceptions can lead to a wide range of equilibrium behavior. When combined with a suitable strategy misperception, *every* correlation misperception is equally viable and the learning channel does not influence the viability of a correlation misperception.

Proposition 3. *Fix any* $r^\bullet > 0$, $\kappa^\bullet \in [0,1]$, $(\sigma_\zeta^\bullet)^2 \geq L$, *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*

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

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

For any correlation misperception, there is some suitably complementary misperception of the residents' strategy inducing the entrants to play the Stackelberg strategy in equilibrium. This misperception takes into account the entrants' equilibrium misinference of the elasticity parameter in the case where they have flexible models. Intuitively, this misperception leads the entrant to believe that the resident's strategy has relatively limited impact on the price, encouraging the use of a more aggressive strategy. We show one can identify a misperception yielding the Stackelberg payoff, the best they can hope for against rational residents. 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, irrespective of the learning channel.

5.1.2 Strategy Misinference

We now consider an alternative way of accommodating entrants who do not know others' strategies: we suppose that entrants draw (possibly wrong) *inferences* about residents' strategy, rather than assuming these misperceptions are fixed as before. We show that this setting delivers a similar message: dropping the assumption that agents correctly know others' strategies in equilibrium prevents us from selecting correlation misperceptions on the basis of equilibrium payoffs.

For entrants with correlation misperception κ_E, a *linear equilibrium with strategy misinference* consists of strategies $\alpha_{R\to R}$, $\alpha_{R\to E}$, $\alpha_{E\to R}$ and beliefs $(\tilde{\kappa}_R, \tilde{r}_R, \tilde{\sigma}^2_{\zeta,R})$, $(\kappa_E, \tilde{\alpha}_R, \tilde{r}_E, \tilde{\sigma}^2_{\zeta,E})$ 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

- $\alpha_{E\to R}$ is an interim-stage best response against $\tilde{\alpha}_R$, given the beliefs $(\kappa_E, \tilde{r}_E, \tilde{\sigma}^2_{\zeta,E})$
- $F^\bullet(\alpha_{E\to R}, \alpha_{R\to E}) = F_{\kappa_E,\tilde{r}_E,\tilde{\sigma}^2_{\zeta,E}}(\alpha_{E\to R}, \tilde{\alpha}_R)$

Like a linear equilibrium with strategy *misperception*, a linear equilibrium with strategy *misinference* features residents who correctly know others' strategies and entrants who have a wrong belief $\tilde{\alpha}_R$ about the residents' stragegy. The key difference is that the entrants now flexibly infer residents' strategy, so $\tilde{\alpha}_R$ is determined in equilibrium. The entrants jointly infer residents' strategy and the persistent parameters of the game ($\tilde{r}_E$ and $\tilde{\sigma}^2_{\zeta,E}$) as to match their observed distribution of market prices, under their dogmatic misperception of signal correlation κ_E.

Proposition 4. *Fix any* $r^\bullet > 0$, $\kappa^\bullet \in [0,1]$, $(\sigma^\bullet_\zeta)^2 \geq L$, *and* $\kappa_E \in [0,1]$. *There exists a linear equilibrium with strategy misinference where the entrants' fitness is weakly lower than the residents' fitness. But there also exists a different linear equilibrium with strategy misinference where the entrants' fitness is strictly higher than the residents' fitness.*

Proposition 4 says there is a multiplicity of equilibria when entrants can draw misinferences about residents' strategy. In fact, for *any* entrant misperception about signal correlation, there exists an equilibrium where the entrants do strictly better than the residents and another equilibrium where they do weakly worse. We do not obtain any selection of which correlation misperceptions are more viable if entrants also make inferences about residents' strategies in equilibrium.

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

We now discuss another line of reasoning that justifies 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. For a fixed environment, we could achieve the same strategic benefits of a flexible model using a suitably chosen dogmatic model. But, flexible models are *strictly* necessary to ensure that the "right" strategic commitments are made *across* different environments. We use similar ideas to show that models with misspecified higher-order beliefs may be more advantageous than first-order misperceptions of the game's persistent parameters. Misspecified higher-order beliefs, unlike first-order misperceptions, allow agents to make beneficial strategic commitments in games while also avoiding losses in non-strategic environments, since these biases only activate in duopolistic markets and become inert in monopolistic markets.

5.2.1 Stability and Invasion with Heterogeneous Markets

We consider 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.

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), (ϕ^m), the entrant model, and the resident model. 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 under the true parameter r^m. We say that the resident model is *resistant to invasion* from the entrant model 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* $\kappa^\bullet \in [0,1]$, $(\sigma_\zeta^\bullet)^2 \geq L$. *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

We now extend the previous elaboration to formally compare the consequences of higher-order misperceptions and first-order misperceptions when agents sometimes act as monopolists. In such monopoly markets, we find that higher-order errors associated with correlation misperception cause no loss of profit and no belief distortion about price elasticity. By contrast, models that encode first-order errors about the slope or intercept of the demand curve *always* cause losses in the monopoly environment. This provides an evolutionary justification for why higher-order misperceptions may be more likely to arise than first-order misperceptions: only the former can strictly improve the agent's profit in duopolistic markets without lowering the agent's profit in monopolistic markets.

Specifically, the *monopoly market* is as follows. Agents are paired to play this game and each chooses a quantity level. But, while the demand state $\omega \sim \mathcal{N}(0, \sigma_\omega^2)$ 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:

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

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, generalizing the linear equilibrium concept by allowing non-linear equilibrium strategies and equilibrium inferences that do not fully explain the observed consequence distribution. Since the two players in the game do not affect each other's payoffs and observations, equilibrium only needs to specify a single strategy for each population, not multiple strategies to be played against different types of opponents. Specifically, 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 we 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 pair of results precisely illustrate the evolutionary basis for misspecification in *higher-order* beliefs over first-order beliefs. We start with misspecifications of higher-order beliefs, showing 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.*

This result comes from the fact that κ has no impact on behavior in the single-agent case. Formally, the conclusion follows from a pair of observations: first, free parameters can be identified from the price distribution as long as the first-order beliefs about the demand state are correct (even if the agent misperceives the opponent's belief about the state). Second, the agent's subjective best response does not depend on their correlation perception, given their profit does not depend on their opponent's action.

We now provide a sharp converse, showing that this conclusion fails if we instead consider models that encode first-order misperceptions of the persistent game parameters: *any* model with a dogmatically wrong r parameter or dogmatically wrong θ parameter leads 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.*

This result mirrors Proposition 6, in particular showing that in this setting, the individually optimal action *cannot* arise if first-order errors are present. An evolutionary advantage of correlation misperceptions 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 payoffs in at least one. Given that deviations from rationality have no evolutionary advantage in decision problems, this is only possible if the bias is an error that only "activates" in strategic situations but not in decision problems. Proposition 6 shows that the correlation misperceptions is an example of such a bias: it strictly increases the entrant's payoffs in the baseline duopoly markets but performs 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, these biases do not "deactivate" in monopoly markets and Proposition 7 implies they always lead to strict profit losses in such settings.

5.3 Assortative Matching

We conclude by discussing an alternative matching process which reverses the selection pressure observed under uniform matching. The solution concept in Definition 4 implicitly assumes matches are drawn uniformly from the population at large; since residents are dominant and entrants are negligible, this means that entrants invariably play against residents. An alternative, discussed in Alger and Weibull (2013), is that entrants only play against other entrants and residents only play against other residents. Alger and Weibull (2019) further discuss how assortative matching may be more realistic for certain population structures. The following definition captures this notion:

Definition 5. A *linear equilibrium with assortative matching* 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 there exists at least one linear equilibrium with assortative matching and $U^\bullet(\alpha_{R\to R}, \alpha_{R\to R}) \geq U^\bullet(\alpha_{E\to E}, \alpha_{E\to E})$ in every such equilibrium, and we say κ_R is *susceptible to invasion with assortative matching* if there exists at least one linear equilibrium with assortative matching and $U^\bullet(\alpha_{R\to R}, \alpha_{R\to R}) < U^\bullet(\alpha_{E\to E}, \alpha_{E\to E})$ in every such equilibrium.

Definition 5 modifies Definition 4 to assume that inferences and payoffs of entrants are determined from equilibrium play against other entrants, instead of against the residents as in Definition 4. Crucially, 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 as for uniform matching:

Proposition 8. *Fix $r^\bullet > 0$ and $(\sigma^\bullet_\zeta)^2 \geq L$. Assume all agents have flexible models. Then, κ_R is susceptible to invasion with assortative 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—and hence more cooperative—behavior. Figure 2 shows a numerical illustration of entrant fitness under assortative matching for different correlation perceptions

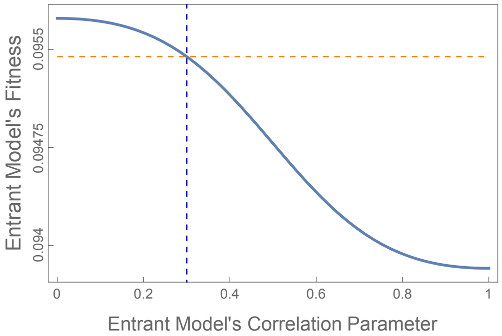

Figure 2: Fitness of flexible entrant models with different correlation perceptions, under assortative matching with a correctly specified resident model $\kappa_R = 0.3$. The true parameters are $\kappa^\bullet = 0.3$, $r^\bullet = 1$, $(\sigma_\zeta^\bullet)^2 = \sigma_\omega^2 = \sigma_\epsilon^2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

κ_E. 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.

Analogous to Proposition 2, the next result shows that the effect of correlation misperception is reversed when the learning channel is shut down: in particular, correlation neglect is always harmful to the entrants. Just as in Proposition 2, we imagine that the entrants initially shared the same (possibly misspecified) flexible model as the residents, and thus they have the same (possibly wrong) belief about price elasticity. We consider the welfare consequence of the entrants switching to a new dogmatic model that freezes their current belief about price elasticity as a fixed parameter but changes their correlation perception.

Proposition 9. *Fix any $r^\bullet > 0$, $\kappa^\bullet \in [0,1]$ and $(\sigma_\zeta^\bullet)^2 \geq L$. Suppose residents have flexible models and perceive κ_R. Suppose entrants have dogmatic models with $\kappa_E \in [0,1]$ and $r_E = r_R$, where r_R is the unique equilibrium belief about elasticity in the resident population induced by κ_R. Then, κ_R is resistant to invasion under assortative matching from entrants with $\kappa_E \leq \kappa_R$. Also, there exists $\epsilon > 0$ so that κ_R is susceptible to invasion under assortative matching from entrants with $\kappa_E \in (\kappa_R, \kappa_R + \epsilon]$.*

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.

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.

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*, 114, 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 (2025): “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 (2025): “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.

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 $\frac{1/\sigma_\epsilon^2}{1/\sigma_\epsilon^2+1/\sigma_\omega^2}$. $\square$

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 $L = \max_{\kappa \in [0,1], 0 \leq \alpha_i, \alpha_{-i} \leq \gamma} \mathrm{Var}_\kappa \left[\omega - \frac{1}{2} r^\bullet \frac{\alpha_i + \alpha_{-i}\psi(\kappa^\bullet)}{\alpha_i + \alpha_{-i}\psi(\kappa)} \alpha_{-i} s_{-i} \mid s_i \right]$. 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 $0 \leq \alpha_i, \alpha_{-i} \leq \gamma$.

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 $(\sigma_\zeta^\bullet)^2 \geq L$, 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 any $\kappa_R, \kappa_E \in [0,1]$, consider a candidate linear equilibrium with strategies $0 \leq \alpha_{R \to R}, \alpha_{E \to R}, \alpha_{R \to E} \leq \gamma$, 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 $r_R := r^\bullet \frac{1 + \psi(\kappa^\bullet)}{1 + \psi(\kappa_R)}$.

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

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

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 that will imply the existence of a unique linear equilibrium for all values of κ_E close enough to κ_R.

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 $H(0) = -(\frac{\gamma}{1+r_R})^2 \frac{1}{2} r^\bullet \psi(\kappa^\bullet)\psi(\kappa_R) + \frac{\gamma^2\psi(\kappa_R)}{1+r_R} = \frac{\psi(\kappa_R)\gamma^2}{1+r_R}(1 - \frac{(1/2)r^\bullet\psi(\kappa^\bullet)}{1+r_R})$; $H(0) > 0$ holds because $1 + r_R > (1/2)r^\bullet\psi(\kappa^\bullet)$, which follows from $r_R = r^\bullet \frac{1+\psi(\kappa^\bullet)}{1+\psi(\kappa_R)} \geq r^\bullet \frac{1+\psi(\kappa^\bullet)}{2}$ and $r^\bullet > 0$. 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 $r_R = r^\bullet\frac{1+\psi(\kappa^\bullet)}{1+\psi(\kappa_R)} \leq 2r^\bullet$. $\square$

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 $\alpha(\kappa_E) = (\alpha_{R\to R}(\kappa_E), \alpha_{R\to E}(\kappa_E), \alpha_{E\to R}(\kappa_E))$.

Next, we show that $\alpha'_{E\to R}$ evaluated at $\kappa_E = \kappa_R$ is strictly positive. That is, locally around $\kappa_E = \kappa_R$, the equilibrium play $\alpha_{E\to R}$ is strictly increasing in κ_E.

Claim A.2. $\alpha'_{E\to R}(\kappa_R) > 0$.

Proof. 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:

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

This expression is proportional to:

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

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

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

But since $r_R(1+\psi(\kappa_R))=r^\bullet(1+\psi(\kappa^\bullet))$, we have that this is proportional to $r_R>0$. Hence this expression is positive.

The coefficient in front of D is:

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

Make the substitution $\gamma=x^*\cdot\left(1+r_R+\frac{1}{2}\psi(\kappa_R)r_R\right)$, and divide by $x^*>0$

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

We show this expression is negative, which is equivalent to showing

$$\begin{aligned}&\frac{r^\bullet r_R}{2(1+r_R)}\psi(\kappa_R)^2\cdot\psi(\kappa^\bullet)+\left(1+r_R+\frac{1}{2}\psi(\kappa_R)r_R\right)\cdot\left(1-\psi(\kappa_R)^2\frac{r_R}{2(1+r_R)}\right)\\ &\qquad<2+2r^\bullet+\left(1-\frac{r_R}{2(1+r_R)}\psi(\kappa_R)\right)\cdot(\psi(\kappa_R)+\frac{1}{2}r^\bullet\psi(\kappa_R)+r^\bullet\psi(\kappa^\bullet))\end{aligned}$$

First, note that:

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

Therefore, we will be done if we can show the following inequality:

$$1 + r_R + \frac{1}{2}\psi(\kappa_R)r_R < 2 + \frac{3}{2}r^\bullet + \frac{1}{2} \cdot (\psi(\kappa_R) + \frac{1}{2}r^\bullet\psi(\kappa_R) + r^\bullet\psi(\kappa^\bullet)),$$

since adding these two inequalities together implies the original claim—indeed, the original inequality emerges if before adding the inequalities together, we multiply the left hand side by $\left(1 - \psi(\kappa_R)^2 \frac{r_R}{2(1+r_R)}\right)$ and, on the right hand side, multiply the term in parentheses by $\left(1 - \frac{r_R}{2(1+r_R)}\psi(\kappa_R)\right)$ instead of 1/2, as the first modification makes the left hand side smaller and the second makes the right hand side bigger. In fact, we can show this inequality even if we subtract 1 from the right hand side. Note that have $r^\bullet(1 + \frac{1}{2}\psi(\kappa^\bullet)) - r_R(1 + \frac{1}{2}\psi(\kappa_R)) = r^\bullet \frac{(\psi(\kappa_R) - \psi(\kappa^\bullet))}{2(1+\psi(\kappa_R))}$. Using this identity and the other simplifications, we thus have that it suffices to show:

$$0 < \frac{r^\bullet}{2} + r^\bullet \frac{(\psi(\kappa_R) - \psi(\kappa^\bullet))}{2(1+\psi(\kappa_R))} + \frac{1}{2} \cdot (\psi(\kappa_R) + \frac{1}{2}r^\bullet\psi(\kappa_R)).$$

However, $\frac{1}{2} + \frac{(\psi(\kappa_R) - \psi(\kappa^\bullet))}{2(1+\psi(\kappa_R))} > 0$, since $\psi(\kappa^\bullet) < 1$. Thus, we have shown that the coefficient in front of D is strictly negative.

Since the term without D is positive and the coefficient in front of D is strictly negative, then $\frac{d\alpha^M}{d\psi(\kappa_E)}$ evaluated at $\kappa_E = \kappa_R$ must be positive. 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. □

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 $\bar{U}_i(\alpha_i) := \mathbb{E}[s_i^2] \cdot (\alpha_i\gamma - \frac{1}{2}r^\bullet\alpha_i^2 - \frac{1}{2}r^\bullet\psi(\kappa^\bullet)\alpha_i\frac{\gamma - \frac{1}{2}r_R\psi(\kappa_R)\alpha_i}{1+r_R} - \frac{1}{2}\alpha_i^2)$. This is a quadratic function in α_i with $\bar{U}_i'' < 0$, so it has a unique maximizer that we will denote as a function of κ_R, namely $\alpha^*_{E\to R}(\kappa_R)$.

The derivative in α_i is, up to a positive multiplicative constant:

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

So we rearrange to get

$$\alpha^*_{E\to R}(\kappa_R) = \frac{\gamma((1+r_R) - \frac{\psi(\kappa^\bullet)r^\bullet}{2})}{1+r_R+r^\bullet(1+r_R - r_R\frac{\psi(\kappa^\bullet)\psi(\kappa_R)}{2})}.$$

Also, $\bar{U_i}'(\alpha_{RR})$ is (up to a positive multiplicative constant):

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

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

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

Note that $r_R\psi(\kappa_R) = r^\bullet(1+\psi(\kappa^\bullet)) - r_R$. Making this substitution, we have:

$$(1+r_R)((r_R + \frac{r^\bullet(1+\psi(\kappa^\bullet)) - r_R}{2}) - r^\bullet(1+\frac{1}{2}\psi(\kappa^\bullet))) + \frac{1}{4}r^\bullet\psi(\kappa^\bullet)(r^\bullet(1+\psi(\kappa^\bullet)) - r_R).$$

which we view as a convex and quadratic function of r_R denoted $g(r_R)$. This function is, up to a positive multiplicative constant, equal to $\bar{U_i}'(\alpha_{RR})$. The function g obtains its minimum value (over all of $\mathbb{R}$) at $\tilde{r} = \frac{1}{4}(-2+2r^\bullet + r^\bullet\psi(\kappa^\bullet)) < r^\bullet$.

In fact, $r_R(1) \geq \tilde{r}$. Otherwise, since $r_R(\kappa_R)$ is strictly decreasing, it would imply that $\tilde{r} = r_R(\tilde{\kappa}_R)$ for some $\tilde{\kappa}_R \in [0,1]$. Using the expression for $\tilde{r}$, we must have $\psi(\tilde{\kappa}_R) = \frac{2+2r^\bullet+3\psi(\kappa^\bullet)r^\bullet}{-2+2r^\bullet+\psi(\kappa^\bullet)r^\bullet}$. Since $\tilde{\kappa}_R \in [0,1]$, we get $\psi(\tilde{\kappa}_R) > 0$, which implies $-2 + 2r^\bullet + \psi(\kappa^\bullet)r^\bullet > 0$. But this would imply that $\psi(\tilde{\kappa}_R) > 1$, which is impossible for any $\tilde{\kappa}_R \in [0,1]$.

To complete the proof, we consider the two cases.

Case 1: $2 + r^\bullet > \psi(\kappa^\bullet)(2 + r^\bullet + 2\psi(\kappa^\bullet)r^\bullet)$.

We show that this implies $\alpha^*_{E\to R}(1) < \alpha_{R\to R}(1)$. In general, $\alpha^*_{E\to R}(\kappa_R) < \alpha_{R\to R}(\kappa_R)$ if and only if

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

For $\kappa_R = 1$, we get $\psi(\kappa_R) = 1$, so that $r_R = \frac{r^\bullet}{2}(1 + \psi(\kappa^\bullet))$. The above inequality becomes:

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

Simplifying and using that $r^\bullet > 0$, this condition reduces to:

$$2 + r^\bullet > \psi(\kappa^\bullet)(2 + r^\bullet + 2\psi(\kappa^\bullet)r^\bullet).$$

Now $\alpha^*_{E\to R}(1) < \alpha_{R\to R}(1)$ implies that $g(r_R(1)) < 0$, since $\bar{U}_i(\alpha_i)$ is concave and quadratic in α_i for every r_R. Next we evaluate $g(r_R(\kappa^\bullet))$. We know $r_R(\kappa^\bullet) = r^\bullet$, so making this substitution we get $g(r_R(\kappa^\bullet)) = \frac{1}{4}(r^\bullet\psi(\kappa^\bullet))^2 > 0$.

Since $r_R(\kappa_R)$ is strictly decreasing in its argument and since g is a convex quadratic function, g changes from taking negative values to taking positive values at most once as κ_R decreases from 1 to 0. So, there exists an interior value $\kappa^* \in (\kappa^\bullet, 1)$ so that $g(r_R(\kappa^*)) = 0$ and in fact $\alpha^*_{E\to R}(\kappa^*) = \alpha_{R\to R}(\kappa^*)$. Against a resident with misperception κ^*_R, the entrant cannot get a higher payoff in linear equilibrium than the residents even if the entrant can commit to any linear strategy. For any $\kappa_R \in [0, \kappa^*)$, $g(r_R(\kappa_R)) > 0$. For any $\kappa_R \in (\kappa^*_R, 1]$, $g(r_R(\kappa_R)) < 0$.

Combining this with Claim A.1 and Claim A.2 proves the result.

Case 2: $2 + r^\bullet \leq \psi(\kappa^\bullet)(2 + r^\bullet + 2\psi(\kappa^\bullet)r^\bullet)$.

We show then $g(r_R(\kappa_R))$ is strictly positive for all $\kappa_R \in [0,1)$ and $g(r_R(1)) \geq 0$. Recall we have $r_R(1) \geq \tilde{r}$. Then, to show $g(r_R(\kappa_R))$ is strictly positive for all $\kappa_R \in [0,1)$, we only need to show that $g(r_R(1)) \geq 0$. But this is precisely when $\alpha^*_{E\to R}(1) \geq \alpha_{R\to R}(1)$, which is equivalent to $2 + r^\bullet \leq \psi(\kappa^\bullet)(2 + r^\bullet + 2\psi(\kappa^\bullet)r^\bullet)$. Combining this with Claim A.1 and Claim A.2 proves the result. In particular, for the case of $\kappa_R = 1$, we know that for all κ_E in a neighborhood of 1, there exist a unique linear equilibrium and the residents have weakly higher payoff unless $\kappa_E > \kappa_R$, which is impossible since $\kappa_E \leq 1 = \kappa_R$. □

A.5 Proof of Proposition 2

Proof. Take L as in Lemma 3. Let r_R be the residents' 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.

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 $\kappa_E \in [0,1]$.

If $2 + r^\bullet > \psi(\kappa^\bullet)(2 + r^\bullet + 2\psi(\kappa^\bullet) r^\bullet)$, then there exists some $\kappa^* \in (\kappa^\bullet, 1)$ such that $\bar{U}_i'(\alpha_{RR}) > 0$ for all $\kappa_R \in [0, \kappa^*)$, $\bar{U}_i'(\alpha_{RR}) = 0$ for $\kappa_R = \kappa^*$, and $\bar{U}_i'(\alpha_{RR}) < 0$ for all $\kappa_R \in (\kappa^*, 1]$. If $2 + r^\bullet \leq \psi(\kappa^\bullet)(2 + r^\bullet + 2\psi(\kappa^\bullet) r^\bullet)$, then $\bar{U}_i'(\alpha_{RR}) > 0$ for every $\kappa_R \in [0,1)$. We have $\bar{U}_i'(\alpha_{RR}) = 0$ for $\kappa_R = 1$ if $2 + r^\bullet = \psi(\kappa^\bullet)(2 + r^\bullet + 2\psi(\kappa^\bullet) r^\bullet)$ and $\bar{U}_i'(\alpha_{RR}) > 0$ for $\kappa_R = 1$ otherwise. Combined with the fact that we always have $\frac{d\alpha_E}{d\kappa_E} < 0$ for every $\kappa_E \in [0,1]$, we have established the proposition.

□

A.6 Proof of Proposition 3

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:

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

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

$\alpha^\bullet_{-i}(\alpha_i)$ 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 $\alpha^\bullet_{-i}(0) = \frac{\gamma}{1+r^\bullet}$, 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.7 Proof of Proposition 4

Proof. First, we construct a linear equilibrium with strategy misinference where the entrants have the same fitness as the residents. Let α° be the symmetric equilibrium under the true parameters. We construct an equilibrium of the form $\alpha_{R\to R} = \alpha_{R\to E} = \alpha_{E\to R} = \alpha^\circ$, $(\tilde{\kappa}_R, \tilde{r}_R, \tilde{\sigma}^2_{\zeta,R}) = (\kappa^\bullet, r^\bullet, (\sigma^\bullet_\zeta)^2)$, with some equilibrium entrant belief about the resident's strategy $\tilde{\alpha}_R$.

For $x \geq 0$, if entrants believe that residents use the linear strategy x, their inference about r will be $\hat{r} = r^\bullet \frac{\alpha^\circ + \alpha^\circ\psi(\kappa^\bullet)}{\alpha^\circ + x\cdot\psi(\kappa_E)}$. Under this inference, their subjective best response to the strategy

x is given by

$$\alpha_i^{BR}(x;\kappa_E;r^\bullet\frac{\alpha^\circ+\alpha^\circ\psi(\kappa^\bullet)}{\alpha^\circ+x\cdot\psi(\kappa_E)})=\frac{\gamma-\frac{1}{2}\psi(\kappa_E)r^\bullet\frac{\alpha^\circ+\alpha^\circ\psi(\kappa^\bullet)}{\alpha^\circ+x\cdot\psi(\kappa_E)}\cdot x}{1+r^\bullet\frac{\alpha^\circ+\alpha^\circ\psi(\kappa^\bullet)}{\alpha^\circ+x\cdot\psi(\kappa_E)}}. \tag{4}$$

Our conjectured equilibrium exists if there is some $x \geq 0$ so that the above expression is equal to α°. From the proof of Proposition 1, we get the closed-form expression

$$\alpha^\circ=\frac{\gamma}{1+\frac{r^\bullet}{2}(2+\psi(\kappa^\bullet))}=\frac{\gamma}{1+r^\bullet(1+\frac{\psi(\kappa^\bullet)}{2})}.$$

Note that Equation (4) is equal to $\frac{\gamma}{1+r^\bullet(1+\psi(\kappa^\bullet))}$ for $x=0$, which is strictly smaller than α°. The limit of Equation (4) as $x\to\infty$ is equal to $\gamma-\frac{1}{2}r^\bullet(\alpha^\circ+\alpha^\circ\psi(\kappa^\bullet))$. We now show that this limit is strictly larger than α°. By simple algebra, this is true if and only if $\alpha^\circ(1+\frac{1}{2}r^\bullet(1+\psi(\kappa^\bullet)))<\gamma$, which is equivalent to $\alpha^\circ<\frac{\gamma}{1+r^\bullet(1/2+\frac{\psi(\kappa^\bullet)}{2})}$. This is true since $\alpha^\circ=\frac{\gamma}{1+r^\bullet(1+\frac{\psi(\kappa^\bullet)}{2})}$. By the intermediate-value theorem, there is some $x\geq 0$ so that Equation (4) is equal to α°.

Next, consider the linear equilibrium with strategy misperception constructed in the proof of Proposition 3, which is self-confirming for the entrants and gives them a fitness strictly higher than that of the residents. This is also a linear equilibrium with strategy misinference, where the entrants' equilibrium inference $\tilde{\alpha}_R$ about the residents' strategy is equal to the dogmatic misperception $\hat{\alpha}_R$ from Proposition 3. □

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. 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). $\square$

A.12 Proof of Proposition 9

Proof. Consider $\alpha_{E\to E}(\kappa)$ 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 $\kappa < \kappa_R$ and strictly smaller when $\kappa > \kappa_R$. From the proof of Proposition 8, 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 $\alpha_{R\to R} > \alpha_{TEAM}$, 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 $\kappa_l < \kappa_R$. Also, there is some $\epsilon > 0$ so that for $\kappa_h \in (\kappa_R, \kappa_R + \epsilon]$, we get $\alpha_{R\to R} > \alpha_{E\to E}(\kappa_h) > \alpha_{TEAM}$, so entrants with $\kappa_E = \kappa_h$ have strictly higher fitness than residents with κ_R, establishing the second claim. $\square$