

MIT Open Access Articles

Justified Communication Equilibrium

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Clark, Daniel and Fudenberg, Drew. 2021. "Justified Communication Equilibrium." American Economic Review, 111 (9).

As Published: 10.1257/AER.20201692

Publisher: American Economic Association

Persistent URL: <https://hdl.handle.net/1721.1/144468>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Justified Communication Equilibrium[†]

By DANIEL CLARK AND DREW FUDENBERG*

Justified communication equilibrium (JCE) is an equilibrium refinement for signaling games with cheap-talk communication. A strategy profile must be a JCE to be a stable outcome of nonequilibrium learning when receivers are initially trusting and senders play many more times than receivers. In the learning model, the counterfactual “speeches” that have been informally used to motivate past refinements are messages that are actually sent. Stable profiles need not be perfect Bayesian equilibria, so JCE sometimes preserves equilibria that existing refinements eliminate. Despite this, it resembles the earlier refinements D1 and NWBR, and it coincides with them in co-monotonic signaling games. (JEL C70, D82, D83, J23, M51)

Cheap-talk communication is available in many of the settings signaling games are intended to model, and signaling games with or without cheap talk can have a great many equilibria. This paper provides a learning foundation for *justified communication equilibrium* (JCE), which is a new equilibrium refinement for signaling games with costly signals and cheap-talk messages. For a given signal and strategy profile, a sender type is *justified* if some conceivable (i.e., undominated) response makes the type weakly prefer to play the signal rather than conform to the strategy profile, and makes all other types weakly prefer to conform. A *justified response* to a signal is a convex combination of best responses to beliefs that assign probability one to the justified types for that signal. JCE requires that for every signal, there is at least one message that induces the receiver to play a justified response.

The restrictions imposed by JCE on off-path play have some of the flavor of commonly used signaling game refinements, such as the Intuitive Criterion (Cho and Kreps 1987) and D1 (Banks and Sobel 1987), but JCE can make very different predictions in economically relevant settings. Unlike those refinements, JCE has a foundation in the theory of learning in games. We provide this foundation by analyzing the limits of steady states in an overlapping generations learning environment where agents are patient, have long expected lifetimes, and the senders on average play many more repetitions of the game than the receivers do. This fits settings where the senders are institutions and the receivers are individuals (or families, clans, etc.),

*Clark: Department of Economics, MIT (email: dgclark@mit.edu); Fudenberg: Department of Economics, MIT (email: drewf@mit.edu). Sylvain Chassang was the coeditor for this article. We thank Glenn Ellison, Kevin He, David Kreps, David K. Levine, Harry Pei, Andy Skrypacz, Joel Sobel, Alex Wolitzky, and two anonymous referees for helpful comments, and NSF grants SES 1643517 and SES 1951056 for financial support.

[†]Go to <https://doi.org/10.1257/aer.20201692> to visit the article page for additional materials and author disclosure statements.

since institutions will typically be involved in many more interactions than individuals. We say that the strategy profiles corresponding to these steady states are *stable*.

We analyze the stable profiles under the assumption that the message space is large enough that, for each signal and subset of sender types, there is a distinct message that claims “I am playing this signal and my type is in this set.” We further assume that receivers are *initially trusting*, which roughly means that the receivers’ prior leads them to trust such messages they have not previously observed to be lies. We view initial trust as a mild and plausible assumption on how receivers respond to messages. Section VC discusses how it relates to past work on the interpretation of communication.

JCE emerges as a necessary condition for stability in our learning model because when senders are long-lived, most of them play a best response to the aggregate play of the receivers. A given signal can only be a best response for justified types, so receivers are very unlikely to encounter a signal being played by a nonjustified type. Initial trust then implies that most receivers will trust a message claiming to be a justified type, and so play a justified response.

Because we formally add cheap talk to the extensive form of the signaling game, our analysis can and does specify how receivers respond to each possible message, including to the “null message” of saying nothing at all, so we can give the first learning foundation for “speeches” of the sort Cho and Kreps (1987) used to motivate the Intuitive Criterion. In particular, these speeches are not counterfactual, but are messages that are actually sent, which lets us determine how receivers respond to them. Thus, we sidestep the “Stiglitz critique” (Cho and Kreps 1987, Rabin and Sobel 1996) of signaling game refinements, which is based on iterated arguments about how players believe their opponent “should” interpret hypothetical deviations, and address the possible complications in adding explicit communication to the signaling game discussed in Fudenberg and Tirole (1991a).

Our results can be seen as both a validation of and a correction to previous signaling game refinements, which are only roughly in line with the implications of nonequilibrium learning. Specifically, none of the traditional equilibrium refinements is a necessary condition for stability in our learning model.¹ Indeed, as shown by Example 3, the stable outcomes of our learning model need not be perfect Bayesian equilibria (Fudenberg and Tirole 1991b), since the response to an off-path signal can be a mixture over pure best responses corresponding to different beliefs that need not itself be a best response to a single belief. For this reason, JCE is not a refinement of perfect Bayesian equilibrium, but instead is a refinement of perfect Bayesian equilibrium with heterogeneous off-path beliefs (PBE-H) (Fudenberg and He 2018).

We explore the relationships of JCE with previous equilibrium refinements later in the paper, but we preview a few results here. As the left-hand box in Figure 1 illustrates, every JCE passes the “Intuitive Criterion test,” and every JCE is a rationality-compatible equilibrium (RCE) (Fudenberg and He 2020), which is the strongest previous equilibrium refinement for signaling games that has a learning foundation. JCE is not nested with D1, but every PBE-H that satisfies NWBR is

¹ Moreover, as far as we know they have not been shown to be necessary in *any* learning model.

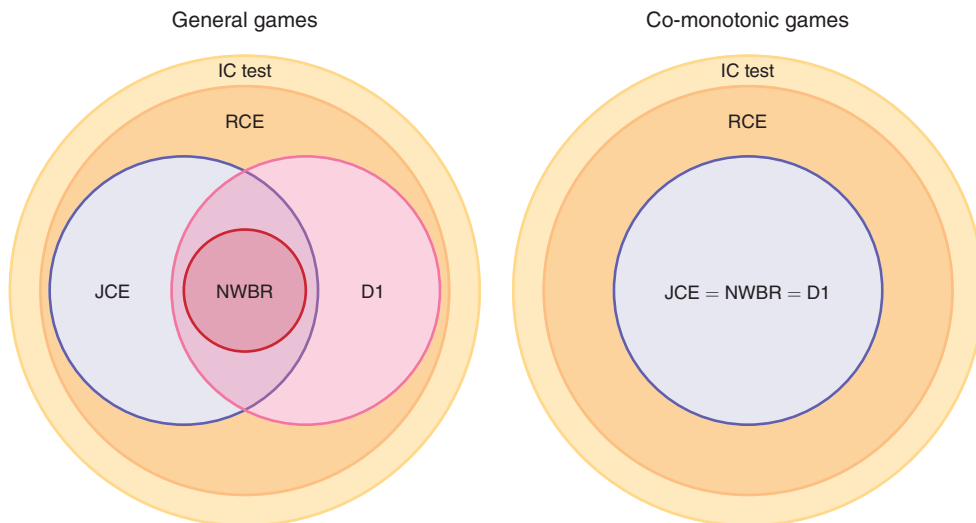


FIGURE 1. JCE AND OTHER RESTRICTIONS OF PBE-H

path equivalent to a JCE.² The right-hand box in Figure 1 depicts the fact that JCE, NWBR, and D1 are essentially equivalent in the special but important class of *co-monotonic* signaling games, which provides a learning justification for selecting the least-cost separating equilibria in many of these games.

I. Preliminaries

A. Signaling Games with Communication

In a signaling game with communication, the sender (player 1) has a type space Θ , a signal space S , and a message space M . The sender observes their type, which is drawn from a full-support distribution $\lambda \in \Delta(\Theta)$, and then chooses a signal $s \in S$ and a message $m \in M$.³ The receiver (player 2) observes the sender's choice of (s, m) , but not the sender's type, then selects their action $a \in A$, after which payoffs are realized. We assume that all of these sets are finite. We denote the set of sender behavior strategies by $\Pi_1 = (\Delta(S \times M))^\Theta$, the set of receiver behavior strategies by $\Pi_2 = (\Delta(A))^{S \times M}$, and let $\Pi = \Pi_1 \times \Pi_2$ be the set of *strategy profiles*.

The utility function of the sender is $u_1 : \Theta \times S \times A \rightarrow \mathbb{R}$ and the utility function of the receiver is $u_2 : \Theta \times S \times A \rightarrow \mathbb{R}$. Each player's utility depends on the sender's signal and type and the receiver's action; neither utility depends on the message of the sender. We will abuse notation slightly and write $u_1(\theta, \pi)$ and $u_2(\pi)$ for the expected payoffs from strategy profile $\pi = (\pi_1, \pi_2)$, as well as $u_1(\theta, s, \alpha)$ for the expected utility of the type θ sender from playing signal s when the receiver responds according to $\alpha \in \Delta(A)$, and $u_2(p, s, \alpha)$ for the expected utility of the

²When we refer to NWBR in this paper we mean "Never a weak best response" in the sense of Cho and Kreps (1987) and Cho and Sobel (1990).

³Throughout, we use $\Delta(\Omega)$ to denote the set of (Borel) probability distributions over a set Ω .

receiver from playing α when the sender plays s and the probability distribution over their type is $p \in \Delta(\Theta)$. Finally, $BR(\theta, s) = \operatorname{argmax}_{a \in A} u_2(\theta, s, a)$ denotes the pure best responses for the receiver to signal s when the sender's type is θ , $BR(p, s) = \operatorname{argmax}_{a \in A} u_2(p, s, a)$ denotes the pure best responses for the receiver to signal s under belief $p \in \Delta(\Theta)$, and $BR(\tilde{\Theta}, s) = \bigcup_{p \in \Delta(\tilde{\Theta})} BR(p, s)$ denotes the pure best responses to signal s for some p with support in $\tilde{\Theta}$.

B. Definition of Justified Communication Equilibrium

The set of actions that are a best response to some belief about θ is $BR(\Theta, s)$. These are the *undominated responses* to s ; the other responses are *conditionally dominated* in the sense of Fudenberg and Tirole (1991a). Thus $\Delta(BR(\Theta, s))$ is the set of receiver mixed actions that assign probability one to undominated responses.

DEFINITION 1 (Fudenberg and He 2018): *Strategy profile π is a PBE-H if*

- (i) For each $\theta \in \Theta$, $u_1(\theta, \pi) = \max_{(s,m) \in S \times M} u_1(\theta, s, \pi_2(\cdot | s, m))$.
- (ii) For each on-path signal-message pair (s, m) , $\pi_2(\cdot | s, m) \in \Delta(BR(p_{(s,m)}, s))$, where $p_{(s,m)}$ is the posterior belief given (s, m) obtained through Bayes' rule.
- (iii) For each off-path signal-message pair (s, m) , $\pi_2(\cdot | s, m) \in \Delta(BR(\Theta, s))$.

Conditions (i) and (ii) of Definition 1 are the conditions for a strategy profile to be a Nash equilibrium. Condition (iii) lets the receiver's response to an off-path signal-message pair (s, m) be a mixture over several actions, each of which is a response to a possibly different belief about the sender's type. Conditions (i)–(iii) together are slightly weaker than perfect Bayesian equilibria (PBE) (Fudenberg and Tirole 1991b). This is because PBE replaces condition (iii) with the requirement that the receiver response to each (s, m) is in the set

$$MBR(\Theta, s) = \{ \alpha \in \Delta(A) : \exists p \in \Delta(\Theta) \text{ such that } u_2(p, s, \alpha) \geq u_2(p, s, a) \forall a \in A \}$$

of mixed best responses to s .⁴ Note, $\Delta(BR(\Theta, s))$ can be strictly larger than $MBR(\Theta, s)$ because it can include mixtures over actions that are not best responses to the same belief.

JCE adds the “justified-response” condition to PBE-H. To define this condition, for each type θ , signal s , and strategy profile π , let

$$\tilde{D}_\theta(s, \pi) = \{ \alpha \in \Delta(BR(\Theta, s)) : u_1(\theta, s, \alpha) > u_1(\theta, \pi) \}.$$

⁴Recall that PBE and sequential equilibrium are equivalent in signaling games.

This is the set of mixtures over undominated receiver responses to s that would make type θ strictly prefer s to their outcome under π .⁵ Let

$$\tilde{D}_\theta^0(s, \pi) = \{\alpha \in \Delta(BR(\Theta, s)) : u_1(\theta, s, \alpha) = u_1(\theta, \pi)\}$$

be the corresponding set for which type θ would be indifferent between s and their outcome under π . For every $s \in S$ and $\pi \in \Pi_1 \times \Pi_2$, let

$$\Theta^\dagger(s, \pi) = \{\theta \in \Theta : \tilde{D}_\theta(s, \pi) \cup \tilde{D}_\theta^0(s, \pi) \not\subseteq \bigcup_{\theta' \neq \theta} \tilde{D}_{\theta'}(s, \pi)\}$$

be the set of types θ for which there is some mixed receiver action $\alpha \in \Delta(BR(\Theta, s))$ that makes θ weakly prefer (s, α) to their outcome under π and no other type θ' strictly prefer (s, α) to their outcome under π .

DEFINITION 2: *The set of justified types for signal s given profile π is*

$$\bar{\Theta}(s, \pi) = \begin{cases} \Theta^\dagger(s, \pi) & \text{if } \Theta^\dagger(s, \pi) \neq \emptyset \\ \Theta & \text{if } \Theta^\dagger(s, \pi) = \emptyset. \end{cases}$$

A justified response to signal s given profile π is a distribution $\alpha \in \Delta(BR(\bar{\Theta}(s, \pi), s))$ that assigns positive probability only to actions that are best responses to beliefs with support in $\bar{\Theta}(s, \pi)$.⁶

Note that in a PBE-H, each type is justified for every signal it plays with positive probability. This is because every signal-message pair the type is willing to play must give them their equilibrium payoff, while no other type can get strictly more than their equilibrium payoff by playing it.

DEFINITION 3: *The strategy profile π is a justified communication equilibrium (JCE) if*

- (i) *it is a PBE-H;*
- (ii) *for each $s \in S$, there is some $m \in M$ such that $\pi_2(\cdot | s, m) \in \Delta(BR(\bar{\Theta}(s, \pi), s))$.*

Every JCE must be a PBE-H. The second condition requires that the receiver's response to each signal is justified for at least one message. Since the equilibrium response to on-path signal-message pairs is justified in any PBE-H, the substance of JCE comes from the requirement that there be a justified response to every off-path signal. As we will see, this conclusion only follows from our learning model when the message space is sufficiently large. However, the definition of JCE applies for

⁵This set is very similar to the set D_θ used by Cho and Kreps (1987) to formulate NWBR; we discuss the differences in Section III.

⁶Appendix AA shows that if π is a PBE-H, $\Theta^\dagger(s, \pi) = \emptyset$ only when s is equilibrium dominated for all types, so how to define $\bar{\Theta}(s, \pi)$ in this case is not important.

any non-null message space, including the case without cheap talk, where the message space is singleton.

C. Hiring a Worker

This subsection presents two stylized examples of a firm (the sender) potentially hiring a worker (the receiver) for a particular job. In both examples, the firm’s signal $s \in \{Hire, Pass\}$ is its choice of whether to hire the worker. The worker’s choice of action $a \in \{e_H, e_M, e_L\}$ represents how hard they work; e_H represents high effort, e_M medium effort, and e_L low effort. The firm has three possible types, $\Theta = \{\theta_H, \theta_M, \theta_L\}$: type θ_H represents high quality, θ_M medium quality, and θ_L low quality.⁷ The payoffs to both parties are normalized to zero when the firm does not hire. The examples differ only in their payoff functions when the firm hires. In the first example JCE rules out an equilibrium that satisfies D1, and in the second JCE preserves an equilibrium ruled out by D1 (and a fortiori ruled out by NWBR). Both of these possibilities can happen in more general settings; our goal here is to illustrate the logic of JCE in a simple and economically sensible setting.

EXAMPLE 1:

| θ_H | e_H | e_M | e_L |
|------------|-------|-------|--------|
| Hire | 16, 2 | 1, 0 | -2, -1 |
| Pass | 0, 0 | 0, 0 | 0, 0 |

| θ_M | e_H | e_M | e_L |
|------------|-------|-------|-------|
| Hire | 8, 0 | 6, 1 | -4, 0 |
| Pass | 0, 0 | 0, 0 | 0, 0 |

| θ_L | e_H | e_M | e_L |
|------------|-------|-------|-------|
| Hire | 4, -1 | 1, 0 | -1, 1 |
| Pass | 0, 0 | 0, 0 | 0, 0 |

In this example, a hired worker wishes to adjust their costly effort with the quality of the firm because the worker gains when the firm does well, and firm quality and worker effort are complements in determining the likelihood of success. Moreover, the return to effort varies with type so much that the intermediate effort level is strictly dominated when probability of the intermediate type θ_M is small.

All firm types have the same ordinal ranking over outcomes, $(Hire, e_H) \succ (Hire, e_M) \succ Pass \succ (Hire, e_L)$, but they do not have the same ranking of outcome distributions. For instance, there are mixtures over e_H and e_L that make θ_H strictly prefer to Hire and θ_L strictly prefer to Pass, while there are mixtures over e_M and e_L that make θ_L strictly prefer to Hire while θ_H strictly prefers to Pass. For motivation, suppose that, relative to the low quality θ_L firm, the high quality θ_H firm can very efficiently capitalize on a worker exerting high effort, but does not benefit much from medium effort, and is harmed by a poor worker exerting low effort. Similarly, there are mixtures over e_M and e_L that make θ_M strictly prefer to Hire and θ_L strictly prefer to Pass, while there are mixtures over e_H and e_L that make θ_L strictly prefer to Hire and θ_M strictly prefer to Pass. This is the case when, relative to the low quality firm, the medium-quality firm gains significantly from medium effort, does not gain much from high effort, and is greatly hurt by low effort.

⁷The conclusions in these examples do not depend on the distribution over types, so we omit λ .

In every JCE there is a positive probability that the worker is hired. To see why, consider a strategy profile π in which $\pi_1(\text{Hire}|\theta) = 0$ for all θ . Observe that when $\alpha \in \Delta(\{e_H, e_M, e_L\})$ satisfies $4\alpha(e_H) + \alpha(e_M) - \alpha(e_L) \geq 0$, either $16\alpha(e_H) + \alpha(e_M) - 2\alpha(e_L) > 0$ or $8\alpha(e_H) + 6\alpha(e_M) - 4\alpha(e_L) > 0$, so either $\tilde{\theta}_H$ or θ_M strictly prefers *Hire* whenever θ_L weakly prefers *Hire*. Thus $\tilde{D}_{\theta_L}(\text{Hire}, \pi) \cup \tilde{D}_{\theta_L}^0(\text{Hire}, \pi) \subseteq \tilde{D}_{\theta_H}(\text{Hire}, \pi) \cup \tilde{D}_{\theta_M}(\text{Hire}, \pi)$, so $\theta_L \notin \Theta^\dagger(\text{Hire}, \pi)$. Moreover, $\Theta^\dagger(\text{Hire}, \pi)$ is not empty, because some effort distributions make θ_H prefer *Hire* and the other types prefer *Pass*, so θ_L is not a justified type.⁸ Since it is optimal for a hired worker to play e_L only when they put positive probability on the firm being θ_L , no justified response can use e_L with positive probability, so all firm types strictly prefer to *Hire*.

Unlike JCE, D1 and weaker equilibrium refinements such as the Intuitive Criterion allow equilibria in which all types *Pass*.⁹ The Intuitive Criterion allows this equilibrium because θ_L would obtain a higher payoff from *Hire* if the worker responds with either e_H or e_M , both of which are undominated. Consequently, the Intuitive Criterion allows the worker respond to *Hire* with e_L , since it is the best response to θ_L . Similarly, D1 allows the worker to respond to *Hire* with e_L , because there is no single type that strictly prefers to *Hire* whenever θ_L weakly prefers to do so. In particular, θ_L strictly prefers to play *Hire* when the worker responds with $(1/7)e_H + (6/7)e_L$, though this makes θ_M strictly prefer to *Pass*. Likewise, θ_L strictly prefers to *Hire* when the worker responds with $(2/5)e_M + (3/5)e_L$, though this makes θ_H strictly prefer to *Pass*.

Example 1 shows that there are some sensible economic environments where JCE makes stronger predictions than D1.¹⁰ The reverse can also be true, as shown in the following example, where JCE allows an outcome that D1 and the stronger NWBR condition rule out. Because JCE, unlike D1 or NWBR, has a learning foundation, this highlights the subtlety of the implications of learning foundations for equilibrium play.

EXAMPLE 2:

| θ_H | e_H | e_M | e_L | θ_M | e_H | e_M | e_L | θ_L | e_H | e_M | e_L |
|-------------|-------|-------|--------|-------------|-------|-------|--------|-------------|-------|-------|-------|
| <i>Hire</i> | 16, 2 | 1, 0 | -5, -3 | <i>Hire</i> | 8, 1 | 6, 0 | -4, -2 | <i>Hire</i> | 4, -8 | 1, 0 | -1, 4 |
| <i>Pass</i> | 0, 0 | 0, 0 | 0, 0 | <i>Pass</i> | 0, 0 | 0, 0 | 0, 0 | <i>Pass</i> | 0, 0 | 0, 0 | 0, 0 |

As before, a hired worker wishes to exert high effort when hired by a high quality firm and low effort when hired by a low quality firm. However, here a hired worker also wishes to exert high effort when hired by a medium quality firm. Moreover, there is no belief over the sender’s type that makes both high and low effort levels best responses due to concavity in the worker’s payoff.

⁸In fact, the set of justified types for *Hire* given π is $\tilde{\Theta}(\text{Hire}, \pi) = \Theta^\dagger(\text{Hire}, \pi) = \{\theta_H, \theta_M\}$.

⁹Section III reviews the formal definitions of the Intuitive Criterion and D1.

¹⁰Section OA.7.2 in the online Appendix provides a qualitatively different example concerning job assignment and corporate culture where JCE is again stronger than D1.

The firm's payoffs are similar to Example 1, except here the payoff of the θ_H firm from a worker exerting low effort is reduced. This guarantees that, whenever a type θ_H or θ_M firm weakly prefers to *Hire* a worker whose effort concentrates on e_H and e_L , type θ_L strictly prefers to *Hire*.

Every type playing *Pass* is both a PBE and a JCE outcome. It is a PBE outcome because all types are deterred from playing *Hire* when the receiver responds with $e_L = BR(\theta_L, \text{Hire})$. Moreover, θ_L is a justified type for *Hire* under a strategy profile π where all types pass since there are effort distributions which make θ_L prefer to play *Hire* and the other types prefer to *Pass*. Thus, e_L is a justified response to *Hire*, so it is a JCE outcome for every type to *Pass*.

However, every type playing *Pass* is not a D1 outcome. This is because no mixed best response to *Hire* puts positive probability on both e_H and e_L . Consequently, every mixed best response that makes θ_L weakly prefer to *Hire*, makes θ_M strictly prefer to do so. Likewise, for type θ_H . The only response to *Hire* allowed by D1 is then $e_H = BR(\theta_M, \text{Hire})$, which deters no type from hiring.

Both Examples 1 and 2 use the setting of a firm hiring a worker, but the point that JCE and D1 are not nested holds more generally, as we explain in Section III.

II. The Learning Model

A. Model Overview

Now we sketch the structure of the learning model we use to provide a foundation for JCE, and then explain why the learning model generates the predictions we saw in the previous examples. (Later subsections provide the remaining details and formal results of the model, as well as some alternative models with the same implications.) The model is an overlapping generations learning environment where time is discrete and doubly infinite, $t \in \{\dots, -2, -1, 0, 1, 2, \dots\}$. For each θ , there is a continuum of agents of mass $\lambda(\theta)$ in the role of a type θ sender, and there is a continuum of agents of mass one in the receiver role. The agents have geometric lifespans: agents in sender roles have continuation probability $\gamma_1 \in [0, 1)$, while agents in the receiver role have continuation probability $\gamma_2 \in [0, 1)$. Each period newborn agents replace the departing agents so the sizes of the various populations are constant, and then agents are anonymously matched into sender-receiver pairs: each sender agent is equally likely to be paired with any of the current receiver agents and vice-versa. In each match, the sender plays a signal s and a message m . The receiver observes the chosen (s, m) and responds with an action a . At the end of each period, both players in a given match observe its outcome, which consists of the type of the sender, the signal and message chosen by the sender, and the action chosen by the receiver.

All agents are rational Bayesians who choose *policies* (maps from past observations to current play) that maximize their expected discounted payoff. At every period t , the *state* of the system is the shares of agents in a given player role with the various possible histories. The state and the optimal policies induce an *aggregate sender strategy* and an *aggregate receiver strategy*, and thus an *update rule* that maps states in period t to states in period $t + 1$. We study this system's steady states, which are the fixed points of the update rule.

Because the receivers observe the type of the sender at the end of each match, neither their continuation probability nor their discount factor impacts their play, and their optimal dynamic programming policy is to simply choose an action that maximizes their expected payoff in the current match. Senders' observations do depend on their play, so their optimal policies incorporate a value for "experimenting" with various signal-message pairs that have the potential to lead to an increase in payoff. The size of the senders' experimentation incentive depends on their continuation probability, their discount factor $\delta \in [0, 1)$, and how much they have already learned: inexperienced senders have more incentive to experiment, and senders cease experimenting when they have enough data. Moreover, different types of sender will choose to experiment in different ways.

We focus on the limits of steady-state play when γ_1 and γ_2 tend to one, so senders and receivers can acquire enough observations to outweigh their prior, and γ_1 tends to one quicker than γ_2 , so that the typical sender plays many more times than the typical receiver. This means that most receivers only ever match with senders who have substantially more experience than them. We also assume that δ goes to one, to ensure that the senders experiment enough to rule out limits that are not Nash equilibria. We call the profiles that correspond to this limit the *stable profiles*.¹¹ This limit provides an idealized version of long-run behavior in settings where the senders are institutions who both have an incentive to experiment and, over time, interact with a large number of individuals in the role of the receivers; one example is firms signaling their knowledge about their productivity, future growth, etc. to potential workers via offers of incentive pay. While workers may interact with a large number of firms over their lifetime, or observe family members and other relations do so, it is unlikely that any given individual will be involved in (or have access to information concerning) as many interactions as the typical large firm.

Preliminary lemmas show that every stable profile must be a PBE-H. The optimality of sender play follows from the fact that patient and long-lived senders eventually stop experimenting and play a best response to the aggregate receiver strategy. The beliefs of long-lived receivers are almost entirely driven by their data as opposed to their priors, ensuring that they respond optimally to on-path signal-message pairs, and because receivers are myopic, their off-path play is always a best response to some beliefs, as PBE-H requires.

Our main result, which shows that all stable profiles are JCE, uses two additional assumptions. First, we assume that the sender message space is sufficiently large that for each signal $s \in S$ and subset of sender types $\tilde{\Theta} \subseteq \Theta$, there is a distinct message $m_{s, \tilde{\Theta}}$ that can be interpreted as "I am playing s and my type is in $\tilde{\Theta}$."¹² We also assume that the receiver "trusts" the message provided that they have not previously encountered a sender with any other type $\theta \notin \tilde{\Theta}$ playing s and

¹¹ As Section VA explains, our results hold under other models of the population structure that also have relatively experienced senders.

¹² The literal content of $m_{s, \tilde{\Theta}}$ need not be "I am playing s and my type is in $\tilde{\Theta}$." Instead, $m_{s, \tilde{\Theta}}$ might be a statement like "I am playing signal s so you should believe my type is in $\tilde{\Theta}$ because ..."

sending $m_{s,\bar{\theta}}$. We discuss these assumptions in more detail in Section IID. With them we prove the following result.

THEOREM 1: *If π is stable, then it is a JCE.*

B. Hiring a Worker, Revisited

We now discuss our learning model in the context of the two “hiring a worker” examples. In particular, we explain why the model rules out the “All *Pass*” outcome in Example 1, where it is consistent with D1 but not JCE, while the model does allow “All *Pass*” in Example 2, where it is consistent with JCE but not D1.

EXAMPLE 1 CONTINUED:

To see why there is no stable profile where all firms pass, recall that with enough experience, firms learn the aggregate effort distribution and exhaust the option value of continued experimentation. Experienced firms then either hire and optimally communicate with workers or *Pass*. If the stable outcome is for all types of firm to *Pass*, it must be that the aggregate effort distribution puts positive probability on effort e_L regardless of how a hiring firm communicates. However, low effort is only optimal for a worker if they assign positive probability to the hiring firm being type θ_L . Initially trusting workers will exert high or low effort when a hiring firm claims to not be type θ_L , unless they have previously experienced deception by type θ_L firms. Since the typical firm has many more interactions over its lifetime than the typical worker, most workers only ever match with experienced firms. Thus, in order for a significant share of workers to experience deception by θ_L firms, θ_L firms must learn that it is optimal to *Hire* and play $m_{In,\{\theta_H,\theta_M\}}$. However, either θ_H or θ_M type firms strictly prefer to *Hire* whenever a θ_L weakly prefers to *Hire*, so one of these types will not *Pass* because it will learn it is strictly optimal to *Hire*.

EXAMPLE 2 CONTINUED:

Online Appendix Section OA.6.1 shows that the outcome where every type plays *Pass* is stable by demonstrating that there are steady-state profiles in which most sufficiently experienced firms play *Pass*, but type θ_L firms experiment with *Hire* much longer than the other types do. This means that the vast majority of workers who have previously been employed were only hired by low quality firms, which leads them to exert effort e_L the next time they are hired. In contrast, initially trusting workers who have not previously experienced employment will exert high effort when first hired by a firm claiming to be of high or medium quality, so the aggregate worker effort distributions will concentrate on e_H and e_L . As observed earlier, under such effort distributions, whenever type θ_H or θ_M weakly prefers to *Hire*, type θ_L strictly prefers to do so. This is what drives type θ_L firms to experiment with *Hire* much more than the other types, which supports the desired steady states.

D1 eliminates the “All *Pass*” outcome while our learning model allows it because D1 only considers receiver mixed best responses. However, in a learning

model, there is no reason that the prevailing aggregate receiver strategy must be a mixed best response, and the steady states described above have aggregate worker responses that put positive probability on both e_H and e_L : inexperienced workers exert effort e_H , while most of the experienced workers learn that it is optimal to exert effort e_L .

C. Details of the Learning Model

We now fill in the remaining details about the learning environment we study, provide formal statements of our assumptions, and prove our main result. We also discuss alternative interpretations of the stable profiles, and related versions of stability that correspond to different ways of passing to the limit. Readers who are more interested in the implications of JCE than its learning foundation can skip ahead to Section III.

At the beginning of their lives, senders have a non-doctrinaire prior $g_1 \in \Delta(\Pi_2)$ over the aggregate receiver behavior strategy π_2 , and receivers have a non-doctrinaire prior $g_2 \in \Delta(\Delta(\Theta \times S \times M))$ over $q \in \Delta(\Theta \times S \times M)$, the prevailing distribution of sender types, signals, and messages.¹³ (To simplify notation, we assume there is a single prior for all agents in a given player role, but all of our results extend to any finite number of priors per role.) Upon observing the outcome of a match, agents update their beliefs in accordance with Bayes' rule, which is always applicable because the priors assign positive probability to any finite sequence of observations. Define $\mathcal{H}_{1,t} = (S \times M \times A)^t$ to be the histories that a sender of age t could have observed, with the convention that $(S \times M \times A)^0 = \emptyset$, and let $\mathcal{H}_1 = \bigcup_{t \in \mathbb{N}} \mathcal{H}_{1,t}$ be the collection of all such histories. Likewise, the relevant pieces of information for the receiver are the type, signal choice, and message choice of the sender. Let $\mathcal{H}_{2,t} = (\Theta \times S \times M)^t$ denote the set of sequences of such triples that a receiver agent with age t could have observed, and let $\mathcal{H}_2 = \bigcup_{t \in \mathbb{N}} \mathcal{H}_{2,t}$ be the collection of all such sequences.

All agents maximize their expected discounted payoff. The receivers use a policy $\mathbf{y} : \mathcal{H}_2 \rightarrow A^{S \times M}$ which maps their histories to pure strategies to maximize

$$\mathbb{E}_{g_2} \left[\sum_{t=0}^{\infty} \sum_{\theta,s,m,a} \gamma^t q(\theta, s, m) u_2(\theta, s, \mathbf{y}(s, m | h_{2,t})) \right].$$

Type θ senders use an optimal policy $\mathbf{x}_\theta^{\delta, \gamma_1} : \mathcal{H}_1 \rightarrow S \times M$ that maps their histories to signal-message pairs to maximize

$$\mathbb{E}_{g_1} \left[\sum_{t=0}^{\infty} \sum_{s,m,a} (\delta \gamma_1)^t \pi_2(a | \mathbf{x}_\theta^{\delta, \gamma_1}(h_{1,t})) u_1(\theta, \mathbf{x}_\theta^{\delta, \gamma_1}(h_{1,t}), a) \right].^{14}$$

¹³Here “non-doctrinaire” means “has a continuous density function that is strictly positive on the interior of the probability simplex.” Since $q(\theta, s, m) = \lambda(\theta) \pi_1(s, m | \theta)$ is the distribution over (θ, s, m) induced by the sender type distribution λ and the aggregate sender behavior strategy π_1 , it would be equivalent to define the receivers' beliefs as elements of $\Delta(\Delta(\Theta) \times \Pi_1)$.

¹⁴Here we slightly abuse notation by having both components of $\mathbf{x}_\theta^{\delta, \gamma_1}$ enter the utility function, though it does not depend on the sender's message.

We will focus on the case where both δ and γ_1 are near one, so the senders have maximal incentives to experiment.¹⁵

At every period t , the *state* of the system, denoted $\mu_t = (\mu_{1,t}, \mu_{2,t}) \in (\Delta(\mathcal{H}_1))^\Theta \times \Delta(\mathcal{H}_2)$, gives the shares of agents in a given player role with the various possible histories. Given μ_t , the profile $\mathbf{x}^{\delta, \gamma_1} = \{\mathbf{x}_\theta^{\delta, \gamma_1}\}_{\theta \in \Theta}$ of sender policies induces a sender behavior strategy $\sigma_1^{\delta, \gamma_1}(\mu_{1,t}) \in \Pi_1$ that we call the *aggregate sender play*. Similarly, the receiver policy \mathbf{y} induces a receiver behavior strategy $\sigma_2(\mu_{2,t}) \in \Pi_2$ that we call the *aggregate receiver play*. We call $\sigma^{\delta, \gamma_1}(\mu_t) = (\sigma_1^{\delta, \gamma_1}(\mu_{1,t}), \sigma_2(\mu_{2,t})) \in \Pi_1 \times \Pi_2$ the *aggregate strategy profile*. (Appendix D gives formal definitions of the mappings $\sigma_1^{\delta, \gamma_1} : (\Delta(\mathcal{H}_1))^\Theta \rightarrow \Pi_1$ and $\sigma_2 : \Delta(\mathcal{H}_2) \rightarrow \Pi_2$, as well as other objects introduced in this subsection.)

A policy profile generates an *update rule* $\mathbf{f}^{\delta, \gamma_1, \gamma_2} : (\Delta(\mathcal{H}_1))^\Theta \times \Delta(\mathcal{H}_2) \rightarrow (\Delta(\mathcal{H}_1))^\Theta \times \Delta(\mathcal{H}_2)$, taking the state in period t to the state in period $t + 1$, a mapping $\mathcal{R}_1^{\delta, \gamma_1} : \Pi_2 \rightarrow \Pi_1$ that describes the limit of the aggregate play of the senders as $t \rightarrow \infty$ when the aggregate receiver play is fixed at π_2 , and a mapping $\mathcal{R}_2^{\gamma_2} : \Pi_1 \rightarrow \Pi_2$ that describes the limit of the aggregate receiver play as $t \rightarrow \infty$ when the aggregate sender play is fixed at π_1 . We refer to the mapping $\mathcal{R}^{\delta, \gamma_1, \gamma_2}(\pi) \equiv (\mathcal{R}_1^{\delta, \gamma_1}(\pi_2), \mathcal{R}_2^{\gamma_2}(\pi_1))$ as the *aggregate response mapping*. Online Appendix Section OA.2.1 verifies that this mapping is continuous.

We study this system's steady states, those μ satisfying $\mathbf{f}^{\delta, \gamma_1, \gamma_2}(\mu) = \mu$. We call the corresponding aggregate strategy profiles the *steady-state profiles*, and denote them by $\Pi^*(g, \delta, \gamma_1, \gamma_2) \subseteq \Pi_1 \times \Pi_2$. As online Appendix Section OA.2.2 shows, these are the fixed points of the aggregate response mapping. Continuity of the aggregate response mapping, along with Brouwer's fixed point theorem, then implies that steady-state profiles always exist.

PROPOSITION 1: $\Pi^*(g, \delta, \gamma_1, \gamma_2)$ consists of the strategy profiles that are fixed points of the aggregate response mapping, and it is nonempty for all $g = (g_1, g_2)$, δ , and γ_1, γ_2 .

We consider the iterated limit $\lim_{\gamma_2 \rightarrow 1} \lim_{\delta \rightarrow 1} \lim_{\gamma_1 \rightarrow 1} \Pi^*(g, \delta, \gamma_1, \gamma_2)$. That is, we focus on strategy profiles that are limits of steady states corresponding to some sequence of parameters δ , γ_1 , and γ_2 satisfying this iterated limit. We will call these the *stable profiles*.¹⁶ A corollary of Proposition 1 is that there are stable strategy profiles.

COROLLARY 1: *Stable strategy profiles exist.*

D. Key Assumptions

Our results about stable profiles use two additional assumptions. First, we assume that the sender message space is sufficiently rich.

¹⁵Recall from Fudenberg and Kreps (1988) and Fudenberg and Levine (1993) that with impatient players, learning need not lead to Nash equilibrium, let alone to refinements of it.

¹⁶Formally, strategy profile π is stable if there is a sequence $\{\gamma_{2,j}\}_{j \in \mathbb{N}} \rightarrow 1$, sequences $\{\delta_{j,k}\}_{j,k \in \mathbb{N}}$ with $\lim_{k \rightarrow \infty} \delta_{j,k} = 1$ for all j , and sequences $\{\gamma_{1,j,k,l}\}_{j,k,l \in \mathbb{N}}$ with $\lim_{l \rightarrow \infty} \gamma_{1,j,k,l} = 1$ for all j, k , such that $\pi = \lim_{j \rightarrow \infty} \lim_{k \rightarrow \infty} \lim_{l \rightarrow \infty} \pi_{j,k,l}$ for some sequence $\pi_{j,k,l} \in \Pi^*(g, \delta_{1,j,k}, \gamma_{1,j,k,l}, \gamma_{2,j})$.

ASSUMPTION 1 (Richness): $|M| \geq 2^{|\Theta|} |S|$.

Assumption 1 requires that the message space is large enough to have a distinct element, $m_{s,\tilde{\Theta}} \in M$, for each signal $s \in S$ and subset of sender types $\tilde{\Theta} \subseteq \Theta$. This allows $m_{s,\tilde{\Theta}}$ to be interpreted as “I am playing s and my type is in $\tilde{\Theta}$.” Our next assumption is that when the sender plays s and sends the message $m_{s,\tilde{\Theta}}$, the receiver “trusts” the message provided that they have not previously encountered a sender with any other type $\theta \notin \tilde{\Theta}$ playing s and sending $m_{s,\tilde{\Theta}}$.

ASSUMPTION 2 (Initial Trust): *For every $s \in S$ and $\tilde{\Theta} \subseteq \Theta$, there is some $m_{s,\tilde{\Theta}} \in M$ such that $\mathbf{y}(s, m_{s,\tilde{\Theta}} | h_2) \in BR(\tilde{\Theta}, s)$ for every $h_2 \in \mathcal{H}_2$ in which $(\theta', s, m_{s,\tilde{\Theta}})$ has not been observed for any $\theta' \notin \tilde{\Theta}$.*

Initial trust says that receivers give the sender the “benefit of the doubt” and act in accordance with certain claims they have not previously seen proved false.¹⁷ It does not require that the receivers are certain that these claims are true, only that they give them a sufficiently high probability of being true. Of course, the receiver may quickly learn to distrust claims that prove to be false, which is why Assumption 2 is only applied to claims for which no direct contradictory evidence exists.¹⁸

We maintain Assumptions 1 and 2 throughout the main text. Section VD discusses alternatives to initial trust that give refinements of PBE-H that are similar to JCE.

Without communication, or with communication but no assumptions on the receivers’ prior beliefs about the meaning of previously unobserved messages, stability offers little predictive power and the theorem is false. In particular, it then allows implausible outcomes, as shown by example in online Appendix Section OA.7.1. In the example, there are two sender types, θ_1 and θ_2 , and two signals, *In* and *Out*. *Out* is strictly dominant for θ_2 , and θ_1 prefers to play *In* if the receiver responds to *In* with the best response to θ_1 , so the reasonable outcome seems to be one where θ_1 plays *In* and θ_2 plays *Out*. Indeed, this is the unique JCE as well as the unique equilibrium outcome that satisfies weaker refinements such as the Intuitive Criterion. However, if the receivers are “initially skeptical” so that when they first witness (In, m) they believe it probably came from θ_2 regardless of m , there are stable profiles in which both types play *Out*. This is because, if very few senders play *In*, the aggregate receiver response to *In* concentrates on the best response to θ_2 , which ensures that almost all senders in the population learn that it is optimal to play *Out*.¹⁹ Intuitively, cheap talk has no effect in “babbling” equilibria where messages are meaningless, and effective communication requires some restrictions on how people interpret messages they have never seen before.

¹⁷Initial trust is similar in spirit to the “believe-unless-refuted” condition of Lipman and Seppi (1995), and is also related to notions of credibility in Rabin (1990), Farrell (1993), and Clark (2020). We discuss these connections in more detail in Section VC.

¹⁸Initial trust implicitly places restrictions on the receivers’ prior g_2 . For simplicity, we state it directly on receiver behavior.

¹⁹The same argument shows that this outcome is also stable when cheap talk is not feasible.

E. Proof of Theorem 1

THEOREM 1: *If π is stable, then it is a JCE.*

To prove this theorem we first show that a stable profile is a PBE-H. Condition (iii) of the definition of PBE-H follows from the fact that the receivers in our model myopically optimize because their observations do not depend on their play. We establish the two other conditions of Definition 1, as well as the additional requirement of JCE given in Definition 3, using three supporting lemmas, whose proofs are in Appendix B.

The following lemma shows that stable profiles satisfy condition (i) of Definition 1.

LEMMA 1: *Suppose that π is stable. Then for each $\theta \in \Theta$, $\pi_1(\cdot|\theta)$ puts support only on those sender signal-message pairs that are optimal for type θ under the receiver behavior strategy π_2 .*

The proof of Lemma 1 shows that for fixed $\gamma_2 \in [0, 1)$, aggregate sender play is optimal given the aggregate receiver play when first $\gamma_1 \rightarrow 1$ and then $\delta \rightarrow 1$. As in Fudenberg and Levine (1993), this holds because each sender type will experiment enough to drive the option value of experimentation to zero, so that aggregate sender play is optimal in the limit. The conclusion of Lemma 1 follows from combining this with the fact that the sender best response correspondence (in the underlying two-player game) has a closed graph.

The next lemma shows that stable profiles satisfy condition (ii) of Definition 1.

LEMMA 2: *Suppose that π is stable. Then for any sender signal-message pair (s, m) that occurs with positive probability under π , $\pi_2(\cdot|s, m)$ puts support only on receiver actions that are best responses to s and the posterior belief induced by λ and $\{\pi_1(s, m|\theta)\}_{\theta \in \Theta}$.*

The proof of Lemma 2 shows that receivers will get enough observations of on-path play for their data to swamp their priors. By the law of large numbers their sample converges to the population distribution with high probability, and since receivers myopically optimize, the lemma follows.

Neither Lemma 1 nor Lemma 2 requires Assumptions 1 or 2. The next lemma does require both assumptions. The lemma shows that, for fixed $s \in S$ and $\tilde{\Theta} \subseteq \Theta$, if every type $\theta \notin \tilde{\Theta}$ strictly prefers their payoff under π to their payoff from playing $(s, m_{s, \tilde{\Theta}})$ (and having the receiver respond with $\pi_2(\cdot|s, m_{s, \tilde{\Theta}})$), then the aggregate receiver response to $(s, m_{s, \tilde{\Theta}})$ must be supported on $BR(\tilde{\Theta}, s)$.

The proof of the lemma, and thus of Theorem 1, fails without Assumption 2, and a fortiori in settings where cheap-talk messages are not available. Moreover, the example in online Appendix Section OA.7.1 shows that without initial trust, there can be stable profiles that are not JCE.

LEMMA 3: *Suppose that π is stable. Fix $s \in S$ and $\tilde{\Theta} \subseteq \Theta$. If $u_1(\theta, s, \pi_2(\cdot|s, m_{s, \tilde{\Theta}})) < u_1(\theta, \pi)$ for all $\theta \notin \tilde{\Theta}$, then $\pi_2(BR(\tilde{\Theta}, s)|s, m_{s, \tilde{\Theta}}) = 1$.*

Here we give some intuition for this result. When $u_1(\theta, s, \pi_2(\cdot | s, m_{s, \tilde{\Theta}})) < u_1(\theta, \pi)$ for all $\theta \notin \tilde{\Theta}$, the proof of Lemma 1 shows that, for fixed γ_2 , the aggregate probability that a type outside of $\tilde{\Theta}$ plays $(s, m_{s, \tilde{\Theta}})$ is small when first $\gamma_1 \rightarrow 1$ and then $\delta \rightarrow 1$. For any fixed receiver continuation probability, the share of receivers in the population who have witnessed a sender with type outside of $\tilde{\Theta}$ play the signal-message pair $(s, m_{s, \tilde{\Theta}})$ becomes arbitrarily small as the aggregate probability of such play by types outside of $\tilde{\Theta}$ approaches zero. Recall that receivers who have never observed a type outside of $\tilde{\Theta}$ play $(s, m_{s, \tilde{\Theta}})$ would respond to $(s, m_{s, \tilde{\Theta}})$ with some action in $BR(\tilde{\Theta}, s)$. Combining these facts, it follows that the share of receivers who play some action in $BR(\tilde{\Theta}, s)$ in response to $(s, m_{s, \tilde{\Theta}})$ becomes arbitrarily close to one in the iterated limit.

PROOF OF THEOREM 1:

Let π be a stable profile. We have already established that π must be a PBE-H. We now show that the justified response condition in Definition 3 holds. Fix $s \in S$. Since π is a PBE-H, $u_1(\theta, s, \pi_2(\cdot | s, m_{s, \tilde{\Theta}(s, \pi)})) \leq u_1(\theta, \pi)$ holds for all $\theta \in \tilde{\Theta}(s, \pi)$. By definition, whenever a receiver response weakly deters all justified types from playing a given signal, it must strictly deter every non-justified type. Thus, $u_1(\theta, s, \pi_2(\cdot | s, m_{s, \tilde{\Theta}(s, \pi)})) < u_1(\theta, \pi)$ for all $\theta \notin \tilde{\Theta}(s, \pi)$. Applying Lemma 3 to $\tilde{\Theta} = \tilde{\Theta}(s, \pi)$ then implies that $\pi_2(\cdot | s, m_{s, \tilde{\Theta}(s, \pi)}) \in \Delta(BR(\tilde{\Theta}(s, \pi), s))$. ■

Theorem 1 shows that only JCE can be stable. Not all JCE are stable, because non-doctrinaire priors prevent receiver agents from ever using weakly dominated strategies, and there can be JCE using weakly dominated receiver strategies. It is difficult to give an exact characterization of stable profiles for general games, because all non-doctrinaire initially trusting priors must be considered to show that a given profile is not stable. Instead, we use direct proofs to show that certain equilibria or classes of equilibria are stable. Proposition C1 in Appendix C gives a partial converse to Theorem 1: it shows that all *uniformly justified* JCE in *strictly monotonic* games are stable for all non-doctrinaire priors, including those that do not satisfy initial trust. We also give direct proofs of stability in Example 2 and most of our other examples. The general approach in these proofs is to modify the aggregate response mapping so that its fixed points coincide with the target strategy profile in the limit, and then show that these fixed points are also fixed points of the true aggregate response mapping.

III. Relation to Other Equilibrium Refinements

We have seen by example that JCE and D1 are not nested. We now study their relationship in more detail, as well as the relationship between JCE and other refinements. As a preliminary step, we show that stable profiles need not be PBE, and a fortiori need not satisfy any refinements of PBE.²⁰ This is the reason that JCE is defined as a refinement of PBE-H.

²⁰The equilibrium refinements in Fudenberg and He (2018) and Fudenberg and He (2020) also relax PBE to PBE-H, but those papers do not show that this relaxation is needed.

EXAMPLE 3:

The type space is $\Theta = \{\theta_1, \theta_2\}$, the signal space is $S = \{In, Out\}$, and the action space is $A = \{a_1, a_2, a_3\}$. The payoffs are given by these tables:

| | | | |
|------------|-------|--------|-------|
| θ_1 | a_1 | a_2 | a_3 |
| <i>In</i> | -2, 1 | 1, 0.1 | 1, -1 |
| <i>Out</i> | 0, 0 | 0, 0 | 0, 0 |

| | | | |
|------------|-------|--------|-------|
| θ_2 | a_1 | a_2 | a_3 |
| <i>In</i> | 1, -1 | 1, 0.1 | -2, 1 |
| <i>Out</i> | 0, 0 | 0, 0 | 0, 0 |

This game does not have a PBE in which both types play *Out*, because θ_1 prefers to play *Out* only if there is positive probability that the receiver responds to *In* with a_1 , while θ_2 prefers to play *Out* only if the receiver’s response to *In* uses a_3 with positive probability; yet there is no mixed best response to *In* where the receiver assigns positive probability to both a_1 and a_3 . Nevertheless, the profile π in which both sender types play *Out* and the receiver always responds to *In* with $(1/2)a_1 + (1/2)a_3$ is a JCE, because both sender types are justified and so $a_1 = BR(\theta_1, In)$ and $a_3 = BR(\theta_2, In)$ are each justified responses.

Moreover, online Appendix Section OA.6.2 shows that both types playing *Out* can be a stable outcome, because there are steady-state profiles in which the aggregate receiver strategy plays a_2 with probability less than one-fourth in response to *In* combined with any message. Under such receiver play, for every message m , it can be optimal for at most one sender type to play (In, m) . Thus, if in the limit the aggregate strategy of type θ_1 plays (In, m) with positive probability, then the aggregate strategy of type θ_2 must play (In, m) with zero probability, and the receivers must learn to respond to (In, m) with $a_1 = BR(\theta_1, In)$. But this response strictly deters type θ_1 from playing (In, m) , and an analogous argument applies for the type θ_2 senders.

Unlike JCE, the Intuitive Criterion (Cho and Kreps 1987), D1 (Banks and Sobel 1987), and NWBR (never a weak best response) (Kohlberg and Mertens 1986, Cho and Kreps 1987) were all formulated as refinements of PBE. However, the procedures they use to restrict out-of-equilibrium beliefs and equilibrium outcomes can be adapted to develop tests for any PBE-H, which lets us more naturally compare the predictions of the modified versions of these refinements with JCE. As we will see, JCE is stronger than the modified version of the Intuitive Criterion. JCE and D1 are not nested, although JCE is nested inside the set of equilibria that satisfy a modified version of D1 we call *co-D1*. JCE and NWBR are particularly similar, and in some sense JCE is an adaptation of NWBR with a learning foundation.

We begin by showing that JCE is stronger than a modified version of the Intuitive Criterion we call the *Intuitive Criterion test*. Let $E(s, \pi) = \{\theta \in \Theta : \max_{a \in BR(\theta, s)} u_1(\theta, s, a) \geq u_1(\theta, \pi)\}$. These are the types for whom s is not equilibrium dominated by profile π in the sense of Cho and Kreps (1987).

DEFINITION 4 (Cho and Kreps 1987): *Strategy profile π passes the Intuitive Criterion test if, for every $s \in S$ and $\theta \in E(s, \pi)$, $\min_{a \in BR(E(s, \pi), s)} u_1(\theta, s, a) \leq u_1(\theta, \pi)$.*

PROPOSITION 2: *If π is a JCE, then π is a PBE-H that passes the Intuitive Criterion test.*

The key step of the proof is to show that in a PBE-H, unless s is equilibrium dominated for every type, s is not equilibrium dominated for any justified type, i.e., $\tilde{\Theta}(s, \pi) \subseteq E(s, \pi)$ when $E(s, \pi) \neq \emptyset$. This implies that if there is a justified response that deters all types from playing s , then the profile passes the Intuitive Criterion test. The proof of Proposition 2 is given in Appendix AA.

To compare JCE with D1 and NWBR, we first develop some notation. For every $\theta \in \Theta$, $s \in S$, and $\pi \in \Pi_1 \times \Pi_2$, let $D_\theta(s, \pi) = \{\alpha \in MBR(\Theta, s) : u_1(\theta, s, \alpha) > u_1(\theta, \pi)\}$ be the set of receiver mixed best responses to s that give type θ strictly more than their equilibrium payoff, and let $D_\theta^0(s, \pi) = \{\alpha \in MBR(\Theta, s) : u_1(\theta, s, \alpha) = u_1(\theta, \pi)\}$ be the mixed best responses that give type θ their equilibrium payoff. These are the analogs of the sets $\tilde{D}_\theta(s, \pi)$ and $\tilde{D}_\theta^0(s, \pi)$ when $\Delta(BR(\Theta, s))$ is replaced by $MBR(\Theta, s)$.

To define D1, let $\Theta^{\ddagger, D1}(s, \pi) = \{\theta \in \Theta : \forall \theta' \neq \theta, D_\theta(s, \pi) \cup D_\theta^0(s, \pi) \not\subseteq D_{\theta'}(s, \pi)\}$, and let $\hat{\Theta}^{D1}(s, \pi) \subseteq \Theta$ be the set of types given by

$$\hat{\Theta}^{D1}(s, \pi) = \begin{cases} \Theta^{\ddagger, D1}(s, \pi) & \text{if } \Theta^{\ddagger, D1}(s, \pi) \neq \emptyset \\ \Theta & \text{if } \Theta^{\ddagger, D1}(s, \pi) = \emptyset. \end{cases}$$

Also, let $MBR(\tilde{\Theta}, s) = \{\alpha \in \Delta(A) : \exists p \in \Delta(\tilde{\Theta}) \text{ such that } u_2(p, s, \alpha) \geq u_2(p, s, a) \forall a \in A\}$ denote the set of mixed best responses to s for beliefs supported on a given $\tilde{\Theta} \subseteq \Theta$.

DEFINITION 5 (Banks and Sobel 1987): *Strategy profile π satisfies D1 if for every $s \in S$, there is an $\alpha \in MBR(\hat{\Theta}^{D1}(s, \pi), s)$ such that $u_1(\theta, s, \alpha) \leq u_1(\theta, \pi)$ for all $\theta \in \Theta$.*

D1 can be stronger than JCE (and rule out some stable profiles) because it only considers receiver mixed best responses, both in finding possible responses to off-path signal-message pairs and in the construction of the sets of sender types to which the receiver must be best responding. As we have seen, however, the larger convex hull of receiver best responses emerges in our learning model rather than the receiver mixed best responses.²¹

To see the difference this makes, for every $s \in S$ and $\pi \in \Pi_1 \times \Pi_2$, let $\Theta^{\dagger, D1}(s, \pi) = \{\theta \in \Theta : \forall \theta' \neq \theta, \tilde{D}_\theta(s, \pi) \cup \tilde{D}_\theta^0(s, \pi) \not\subseteq \tilde{D}_{\theta'}(s, \pi)\}$ be the set of types θ where, for every $\theta' \neq \theta$, there is some mixed receiver action $\alpha \in \Delta(BR(\Theta, s))$ that makes θ weakly prefer (s, α) to their equilibrium outcome and θ' weakly prefer their equilibrium outcome to (s, α) . Let $\tilde{\Theta}^{D1}(s, \pi) \subseteq \Theta$ be the set

$$\tilde{\Theta}^{D1}(s, \pi) = \begin{cases} \Theta^{\dagger, D1}(s, \pi) & \text{if } \Theta^{\dagger, D1}(s, \pi) \neq \emptyset \\ \Theta & \text{if } \Theta^{\dagger, D1}(s, \pi) = \emptyset. \end{cases}$$

²¹Fudenberg and Kreps (1988) and Sobel, Stole, and Zapater (1990) recognized that the convex hull of best responses is more natural in a learning setting, but neither paper showed that restricting attention to the receiver mixed best responses rules out a profile that is stable in a learning model.

DEFINITION 6: A PBE-H π is co-D1 if for every $s \in S$, there is an $\alpha \in \Delta(BR(\tilde{\Theta}^{D1}(s, \pi), s))$ such that $u_1(\theta, s, \alpha) \leq u_1(\theta, \pi)$ for all $\theta \in \Theta$.

PROPOSITION 3: If π is a JCE, then π is a PBE-H that is co-D1.

Co-D1 is more permissive than JCE because it strikes fewer types. Appendix AB gives the proof of Proposition 3, which shows that $\tilde{\Theta}(s, \pi) \subseteq \tilde{\Theta}^{D1}(s, \pi)$ for all s ; Example 1 shows that the inclusion is sometimes strict.

To define NWBR, let $\Theta^\dagger(s, \pi) = \{\theta \in \Theta : D_\theta^0(s, \pi) \not\subseteq \bigcup_{\theta' \neq \theta} D_{\theta'}(s, \pi)\}$, which are the θ for which there is a mixed receiver best response $\alpha \in MBR(\Theta, s)$ that makes θ indifferent between (s, α) and their equilibrium outcome and every other type weakly prefer their equilibrium outcome to (s, α) . Let $\hat{\Theta}(s, \pi) \subseteq \Theta$ be the set

$$\hat{\Theta}(s, \pi) = \begin{cases} \Theta^\dagger(s, \pi) & \text{if } \Theta^\dagger(s, \pi) \neq \emptyset \\ \Theta & \text{if } \Theta^\dagger(s, \pi) = \emptyset. \end{cases}$$

DEFINITION 7 (Kohlberg and Mertens 1986, Cho and Kreps 1987): Strategy profile π satisfies NWBR if, for every $s \in S$, there is some $\alpha \in MBR(\hat{\Theta}(s, \pi), s)$ such that $u_1(\theta, \alpha) \leq u_1(\theta, \pi)$ for all $\theta \in \Theta$.

Up to path equivalence, JCE selects the same profiles as NWBR would if the mixed best responses $MBR(\tilde{\Theta}, s)$ were replaced with the convex hulls of best responses $\Delta(BR(\tilde{\Theta}, s))$. Indeed, as shown in online Appendix Section OA.1, it would be equivalent to define JCE by setting $\Theta^\dagger(s, \pi) = \{\theta \in \Theta : \tilde{D}_\theta^0(s, \pi) \not\subseteq \bigcup_{\theta' \neq \theta} \tilde{D}_{\theta'}(s, \pi)\}$, rather than $\Theta^\dagger(s, \pi) = \{\theta \in \Theta : \tilde{D}_\theta(s, \pi) \cup \tilde{D}_\theta^0(s, \pi) \not\subseteq \bigcup_{\theta' \neq \theta} \tilde{D}_{\theta'}(s, \pi)\}$. Thus JCE modifies NWBR in much the same way that co-D1 modifies D1, so NWBR is a stronger refinement than JCE.

PROPOSITION 4: Any PBE-H that satisfies NWBR is path equivalent to a PBE that is a JCE.²²

A PBE-H π that satisfies NWBR is not necessarily a PBE, since the receiver's response to off-path play need not be a best reply to any single belief over the sender's type. However, every such profile is path equivalent to a PBE, since the receiver's response to a given off-path (s, m) can always be replaced by some $\alpha \in MBR(\hat{\Theta}(s, \pi), s)$ that deters the sender types from playing it. Appendix AC completes the proof of Proposition 4 by showing that an "NWBR type" is always a justified type. That is, for a given signal and PBE-H, $\hat{\Theta}(s, \pi) \subseteq \tilde{\Theta}(s, \pi)$ for all s and all PBE-H π .

The converse of Proposition 4 is in general false, as shown earlier by Example 2. However, there are important settings in which NWBR and JCE are path equivalent. One is when there are at most two undominated receiver responses to each

²²Path equivalence is needed in this statement because, unlike JCE, NWBR does not impose requirements about the receiver's actual responses to off-path play.

signal, because then mixed best responses and convex hulls of best responses are the same. We now explore a different class of games where this equivalence holds.

IV. Co-monotonic Signaling Games

This section highlights an important class of commonly studied games in which JCE, D1, and NWBR are path equivalent, so the learning foundation for JCE applies to D1 and NWBR as well. In these *co-monotonic* signaling games, all sender types share the same preference over mixtures over $BR(\Theta, s)$.

DEFINITION 8: *A signaling game is co-monotonic if, for all $\theta, \theta' \in \Theta$, $s \in S$, and $\alpha, \alpha' \in \Delta(BR(\Theta, s))$, $u_1(\theta, s, \alpha) \geq u_1(\theta, s, \alpha')$ if and only if $u_1(\theta', s, \alpha) \geq u_1(\theta', s, \alpha')$.*

This is a subset of the *monotonic* signaling games studied in Cho and Sobel (1990), where the sender types are required to share the same preference only over the receiver mixed best responses $MBR(\Theta, s)$ rather than the convex hull of those responses.

A sufficient condition for a signaling game to be co-monotonic is that there be functions $v : S \times A \rightarrow \mathbb{R}$, $\omega : \Theta \times S \rightarrow \mathbb{R}_{++}$, and $\psi : \Theta \times S \rightarrow \mathbb{R}$ such that $u_1(\theta, s, a) = \omega(\theta, s)v(s, a) + \psi(\theta, s)$ for all $\theta \in \Theta$, $s \in S$, and $a \in A$. Many games, including the following simple economic example, satisfy this condition.

EXAMPLE 4:

Like Examples 1 and 2, this example concerns a firm hiring a worker, except here the firm offers incentive pay to their prospective employee. The firm is better informed about the productivity of the worker's effort; its information is represented by its type $\theta \in \Theta = \{1, 2, 3\}$, with each type equally likely. The firm's signal $s = (s_1, s_2) \in \{0, 1/4, 1/2, 3/4, 1\} \times \{0, 1, 2, \dots, 100\}$ consists of a share of profits s_1 and a base wage s_2 which the worker is offered, and the action $a \in \{0, 5, 10, \dots, 60\}$ represents the worker's choice of effort level. The expected profit given the firm's type θ and the worker's effort a is θa . Thus, the payoffs to the sender and receiver are $u_1(\theta, s, a) = \theta(1 - s_1)a - s_2$ and $u_2(\theta, s, a) = \theta s_1 a + s_2 - a^2/40$, which satisfy the sufficient condition for co-monotonic signaling games given above. Section OA.6.3 in the online Appendix shows that JCE selects equilibria that approximate the least-cost separating equilibrium of this game.

We now explore JCE's relationship with other refinements in co-monotonic games. Co-monotonicity implies that, for all s , any mixture over receiver best responses $\alpha \in \Delta(BR(\Theta, s))$ has a corresponding receiver mixed best response $\alpha' \in MBR(\Theta, s)$ such that $u_1(\theta, s, \alpha) = u_1(\theta, s, \alpha')$ for all θ . This ensures that $\bar{\Theta}(s, \pi) = \hat{\Theta}(s, \pi)$ for every PBE-H π .

LEMMA 4: *In a co-monotonic signaling game, $\bar{\Theta}(s, \pi) = \hat{\Theta}(s, \pi)$ for all $s \in S$ and PBE-H $\pi \in \Pi$.*

The proof of Lemma 4 is in Appendix AC.

In co-monotonic games, all types agree about which receiver best responses are least desirable. Combining this with Lemma 4 shows that JCE and NWBR (Definition 7) select the same profiles up to path equivalence. JCE thus provides a learning foundation for the predictions of NWBR in the class of co-monotonic games.

PROPOSITION 5: *In a co-monotonic signaling game, every JCE is a PBE-H that satisfies NWBR, and every PBE-H that satisfies NWBR is path equivalent to a JCE.*

PROOF:

Suppose that π is a PBE-H that satisfies NWBR. Then, by Proposition 4, π is path equivalent to a JCE.

If π is a JCE, it is a PBE-H. Moreover, for every $s \in S$, there is some $\alpha_s \in \Delta(BR(\bar{\Theta}(s, \pi), s))$ such that $u_1(\theta, s, \alpha_s) \leq u_1(\theta, \pi)$ for all $\theta \in \Theta$. Because the game is co-monotonic, there exists $a_s \in BR(\bar{\Theta}(s, \pi), s)$ such that $a_s \in \arg \min_{a \in BR(\bar{\Theta}(s, \pi), s)} u_1(\theta, s, a)$ for all $\theta \in \Theta$, so $u_1(\theta, s, a_s) \leq u_1(\theta, \pi)$ for all $\theta \in \Theta$. Since the game is co-monotonic, Lemma 4 implies that $a_s \in BR(\bar{\Theta}(s, \pi), s)$, so π is a PBE-H that satisfies NWBR. ■

Combining Proposition 5 with the observation that every PBE-H that satisfies NWBR is path equivalent to a PBE shows that in co-monotonic signaling games, every JCE is path equivalent to a PBE that satisfies NWBR. Moreover, as shown by Cho and Sobel (1990), NWBR and D1 coincide in monotonic games, so JCE is also path equivalent to D1 in co-monotonic games.

COROLLARY 2: *In a co-monotonic signaling game, every JCE is path equivalent to a PBE that satisfies NWBR and D1, and every PBE-H that satisfies NWBR or D1 is path equivalent to a JCE.*

Thus, JCE provides a learning foundation for restricting attention to D1 equilibria in co-monotonic games, as in, e.g., Nachman and Noe (1994); DeMarzo and Duffie (1999); and DeMarzo, Kremer, and Skrzypacz (2005).²³

In various co-monotonic games, such as that of DeMarzo and Duffie (1999), JCE selects the least-cost separating equilibrium outcome, often called the ‘‘Riley outcome’’ (Riley 1979). Moreover, Cho and Sobel (1990) showed that NWBR selects the Riley outcome in a class of monotonic games with a continuum of actions. The definition of JCE can be applied as is to signaling games with infinite actions, and the equivalence of JCE and NWBR in Proposition 5 continues to hold in all co-monotonic signaling games. Thus, JCE selects the Riley outcome in all co-monotonic games that satisfy the additional assumptions of Cho and Sobel (1990) and, by a closed graph argument, also only selects equilibria that are close to the Riley outcome when the action space is a sufficiently fine finite grid, as in Example 4.²⁴

²³Technically, the game analyzed in DeMarzo, Kremer, and Skrzypacz (2005) is not a traditional signaling game because of the presence of multiple senders, but this distinction is not important.

²⁴As noted by, e.g., Fudenberg and Tirole (1991a), it may seem odd that adding a type with a small probability ε can make a large change in the Riley outcome. Stability tracks this change in the Riley outcome because we hold the prior fixed as we take the iterated limit.

V. Discussion

A. Alternate Models

The key to our analysis is that we consider a limit where most senders and receivers have substantial experience, but typical senders have significantly more experience, so that most receivers never encounter inexperienced senders. Because it is inexperienced senders who are the most likely to “experiment” with signal-message pairs that depart from the limit strategy profile, most receivers have little experience with off-path play by the senders, which facilitates the analysis of the stable profiles.

We can obtain this situation with many different specifications of the populations of agents and how they interact. For example, suppose that senders and receivers have geometric lifetimes with common continuation probability $\gamma \in [0, 1)$, so that they all have expected lifetime $T = 1/(1 - \gamma)$. Every period, each sender is matched with a receiver, but each receiver only gets matched with some i.i.d. probability $p \in (0, 1)$.²⁵ A given receiver is expected to have $N_2 = pT$ matches over their lifetime, while a sender is expected to have $N_1 = T$ matches. For every steady state in this alternate model, there is a corresponding steady state in our main model with the same aggregate strategy profile when the receiver’s continuation probability is $\tilde{\gamma}_2 = (1 - 1/T)N_2/(1 + (1 - 1/T)N_2)$, which we demonstrate in online Appendix Section OA.10. Since $\tilde{\gamma}_2 \rightarrow N_2/(1 + N_2) \in [0, 1)$ as $T \rightarrow \infty$ for any fixed $N_2 \in \mathbb{R}_+$ and $N_2/(1 + N_2) \rightarrow 1$ as $N_2 \rightarrow \infty$, the iterated limit where first $T \rightarrow \infty$ (so that both sender and receiver agents become long-lived) then $\delta \rightarrow 1$ (so that sender agents become patient) then $N_2 \rightarrow \infty$ (so that receiver agents become experienced) generates precisely the same predictions as our notion of stability.

Moreover, we can also obtain the same set of stable profiles in models where agents do not have geometric lifetimes: to illustrate, suppose that agents have deterministic lifetimes, and that sender agents are matched every period, while receiver agents are matched every K periods during their life. Suppose that sender agents are involved in N_1 matches over the course of their lifetime, while receiver agents are involved in N_2 . Focusing on the profiles that emerge in the limit where first $N_1 \rightarrow \infty$ then $\delta \rightarrow 1$ then $N_2 \rightarrow \infty$ generates exactly the same predictions as stability in the geometric lifetime models. Thus, the unequal lifetimes of our baseline model are simply a modeling convenience, and not an essential feature.

However, we do need some sort of asymmetry in the interaction structure to derive our results. When both populations have the same expected number of interactions, we have not been able to derive interesting restrictions on the stable profiles. When the receivers have many more interactions (as would be the case in our baseline model in the limit where first $\gamma_2 \rightarrow 1$ then $\gamma_1 \rightarrow 1$), all stable profiles must be PBE and not only PBE-H, but we do not know whether stability has additional implications.

²⁵ Correspondingly set the population mass of senders to be p times that of the receivers.

B. Orders of Limits

We require senders to become long lived ($\gamma_1 \rightarrow 1$) before they become patient ($\delta \rightarrow 1$) so that near the limit very few senders will choose to experiment. (This is why we need the model to have a discount factor parameter.) As in Fudenberg and Levine (1993, 2006) and Fudenberg and He (2018, 2020, 2021), it seems difficult to establish in general what happens when a player's patience level δ goes to one before their lifetime becomes long, as in this case we do not know how to show that most players stop experimenting. The order with which γ_2 and δ go to one is not crucial; we specify that δ converges to one before γ_2 because it affords slightly cleaner results and simpler proofs. All profiles that we prove are stable in our examples would also be stable under a more general version of the iterated limit where first $\gamma_1 \rightarrow 1$ and then $(\delta, \gamma_2) \rightarrow (1, 1)$. Moreover, online Appendix Section OA.9 shows that Theorem 1's conclusion applies under this general limit to all stable profiles satisfying an additional condition, such as on-path strict incentives for the receiver.

C. Related Work

Fudenberg and Kreps (1988) introduced the analysis of nonequilibrium learning in extensive-form games, and announced a program of deriving equilibrium refinements from learning foundations, but did not provide details. Our steady-state formulation is in the spirit of Fudenberg and Levine (1993). Fudenberg and Levine (1993) and Fudenberg and Kreps (1994) provided conditions for rational players to do enough experimentation to rule out non-Nash outcomes.²⁶ Fudenberg and Levine (2006) used a steady-state learning model to study equilibrium refinements in a class of games of perfect information, and showed that all "subgame-confirmed" equilibria are stable.

In signaling games without cheap talk, Fudenberg and He (2018) analyzed the steady states of a model where senders and receivers have identically distributed geometric lifetimes. It assumed that the senders' prior beliefs over the aggregate receiver responses are independent across signals, so that the senders' optimal policy is given by the Gittins index (Gittins 1979), and used this to derive restrictions on equilibria. Fudenberg and He (2020) extended Fudenberg and He (2018) by supposing that the senders assign probability zero to receivers playing conditionally dominated actions, and gave a learning foundation for rationality-compatible equilibrium (RCE). If we treat the signal-message pair (s, m) as a signal when evaluating the conditions of RCE, then RCE collapses to the Intuitive Criterion when the message space is not singleton because messages have no effect on payoffs. If we instead compare RCE to JCE in a game with a singleton message space, RCE is again weaker. For example, the "All Pass" outcome of Example 1 is consistent with

²⁶ Kalai and Lehrer (1993), Lehrer and Solan (2007), Esponda (2013), and Battigalli et al. (2019) studied rational learning without assuming that agents are patient. Battigalli (1987); Rubinstein and Wolinsky (1994); Dekel, Fudenberg, and Levine (1999); Esponda (2013); Battigalli et al. (2015); and Fudenberg and Kamada (2015, 2018), among other papers, studied equilibrium concepts motivated by rational learning without analyzing an explicit learning model, and, e.g., Binmore and Samuelson (1999), Nöldeke and Samuelson (1993), Hart (2002), Jehiel and Samet (2005) studied evolutionary or boundedly rational learning dynamics in extensive form games.

RCE but not JCE.²⁷ Moreover, online Appendix Section OA.3 shows that every JCE is an RCE, because types that are “less compatible” with a given signal in the sense of RCE can never be justified. This paper obtains a stronger refinement than RCE without assuming independent priors by explicitly modeling cheap-talk messages and combining this with the assumptions of initially trusting receivers and relatively long-lived senders.

We view initial trust as a plausible and appealingly simple assumption. It has a similar form to the “believe-unless-refuted” condition of Lipman and Seppi (1995), which is an equilibrium refinement for signaling games with multiple receivers and partial provability. There, each receiver can learn from refutations provided by other receivers. Initial trust is also related to the restrictions imposed by Rabin (1990), Farrell (1993), and Clark (2020) on how receivers respond to “credible” messages in signaling games.²⁸ In these papers, common knowledge of the equilibrium to be played figures heavily in determining the credibility of messages; such restrictions do not fit with our model of nonequilibrium learning. Moreover, deriving restrictions on equilibria from a learning model yields more insight than imposing the restrictions directly.

D. Extensions

We can obtain similar solution concepts by replacing initial trust with alternative assumptions. For example, if receivers know the payoff functions of the senders, as in Fudenberg and He (2020), then receivers who are long-lived may feel that they have acquired a good sense of each sender type’s equilibrium payoff. In online Appendix Section OA.8.1, we discuss a weakened version of initial trust which only requires receivers to trust previously unencountered claims if they are consistent with the receiver’s evaluation of the senders’ incentives. Any stable profile under this assumption must satisfy a refinement that is similar to, but weaker than, JCE. Online Appendix Section OA.8.2 shows that we can capture an iterated procedure similar to that of *divinity* and *universal divinity* (Banks and Sobel 1987) by strengthening initial trust: if the only types who have lied about being in $\tilde{\Theta}$ are elements of $\tilde{\Theta}'$, then the receiver responds to a claim of $\tilde{\Theta}$ as if the sender’s type belongs to $\tilde{\Theta} \cup \tilde{\Theta}'$.

An extensive experimental literature shows that a nontrivial share of experimental subjects tell the truth even when this earns less compensation, and thus behave as if they face a cost of lying. (See the papers surveyed in Abeler, Nosenzo, and Raymond 2019). Kartik, Ottaviani, and Squintani (2007) and Kartik (2009) incorporated messages with such lying costs into models of strategic communication. Online Appendix Section OA.8.3 discusses how our analysis can be extended to signaling games with costly lying. Intuitively, lying costs make it less appealing for a nonjustified type to falsely represent itself as justified.

²⁷RCE also permits equilibria ruled out by JCE and D1 in co-monotonic games like Example 4.

²⁸Rabin (1990) and Farrell (1993) only analyzed cheap-talk games, but their refinements can be extended to games where the sender also has costly signals. Matthews, Okuno-Fujiwara, and Postlewaite (1991); Blume and Sobel (1995); Zapater (1997); Olszewski (2006); Chen, Kartik, and Sobel (2008); and Gordon et al. (2021) also studied refinements in cheap-talk games.

Finally, JCE has no cutting power in games where the sender's only actions are cheap-talk messages. Developing learning foundations for refinements in these games is a promising area for future research, and could lead to learning-based refinements for settings with cheap talk and multiple audiences, as in Goltsman and Pavlov (2011).

VI. Conclusion

Adding cheap-talk communication to signaling games let us provide a learning-theoretic foundation for the concept of JCE. We recovered some of the intuitions that underlie traditional equilibrium refinements for signaling games, whose predictions were by and large sensible in the games where they were used. We also confirmed that some of the worries in the literature about the details of these refinements were well-founded, and pointed out how those refinements need to be modified to accord with the implications of nonequilibrium learning.²⁹

Of course, there are multiple ways to formulate models of nonequilibrium learning, just as there are many definitions of forward induction, and several variants of the Kohlberg and Mertens (1986) axioms. In our opinion, it is easier to judge the plausibility of assumptions on learning models than of axiomatic conditions on equilibrium concepts, especially axioms that are imposed without any reference to how equilibrium play might arise. For this reason, our work makes a valuable contribution even in settings such as co-monotonic signaling games, where the predictions of JCE coincide with those of past work. Outside of those cases, not only does JCE have the benefit of a learning foundation, it is also easier to compute, which may make it more appealing to use.

APPENDIX A. OTHER REFINEMENTS

A. Intuitive Criterion

LEMMA A1: *If π is a JCE, then, for every $s \in S$, either*

$$(i) \Theta^\dagger(s, \pi) \neq \emptyset, \text{ or}$$

$$(ii) u_1(\theta, s, a) < u_1(\theta, \pi) \text{ for all } \theta \in \Theta \text{ and } a \in BR(\Theta, s).$$

PROOF:

Let π be a JCE. Fix $s \in S$ and suppose that $\Theta^\dagger(s, \pi) = \emptyset$. Let $\mathcal{A}_- = \{\alpha \in \Delta(BR(\Theta, s)) : u_1(\theta, s, \alpha) < u_1(\theta, \pi) \forall \theta \in \Theta\}$ be the set of mixtures over receiver best responses that make playing s strictly worse for every type than their outcome under π . Similarly, let $\mathcal{A}_+ = \{\alpha \in \Delta(BR(\Theta, s)) : \exists \theta \in \Theta \text{ such that } u_1(\theta, s, \alpha) > u_1(\theta, \pi)\}$ be the set of mixtures over receiver best responses that make some type strictly better off by playing s than under π . Note, \mathcal{A}_- and \mathcal{A}_+ are disjoint open subsets of $\Delta(BR(\Theta, s))$, and $\mathcal{A}_- \cup \mathcal{A}_+ = \Delta(BR(\Theta, s))$ since $\Theta^\dagger(s, \pi) = \emptyset$.

²⁹The similarity of stable profiles and those that survive NWBR does not hold in general outside of signaling games; see footnote 10 of Fudenberg and Levine (1993).

As $\Delta(BR(\Theta, s))$ is connected, either $\Delta(BR(\Theta, s)) = \mathcal{A}_-$ or $\Delta(BR(\Theta, s)) = \mathcal{A}_+$. Notice, $\Delta(BR(\Theta, s)) = \mathcal{A}_+$ is not possible when π is a JCE since then, for every $\alpha \in \Delta(BR(\bar{\Theta}(s, \pi), s))$, there is a θ such that $u_1(\theta, s, \alpha) > u_1(\theta, \pi)$. Thus $\Delta(BR(\Theta, s)) = \mathcal{A}_-$, so $u_1(\theta, s, a) < u_1(\theta, \pi)$ for all $a \in BR(\Theta, s)$. ■

PROOF OF PROPOSITION 2:

If $E(s, \pi) \neq \emptyset$, there is some θ and $\alpha \in BR(\Theta, s)$ such that $u_1(\theta, s, a) \geq u_1(\theta, \pi)$. By Lemma A1, $\Theta^\dagger(s, \pi) \neq \emptyset$, so $\bar{\Theta}(s, \pi) = \Theta^\dagger(s, \pi)$. Moreover, $\bar{\Theta}(s, \pi) \subseteq E(s, \pi)$, because $\max_{a \in BR(\Theta, s)} u_1(\theta, s, a) < u_1(\theta, \pi)$ implies $\tilde{D}_\theta(s, \pi) \cup \tilde{D}_\theta^0(s, \pi) = \emptyset \subseteq \tilde{D}_{\theta'}$ for any $\theta' \in \Theta$. Thus, $BR(\bar{\Theta}(s, \pi), s) \subseteq BR(E(s, \pi), s)$. Hence, for all $\theta \in \Theta$, $\min_{a \in BR(E(s, \pi), s)} u_1(\theta, s, a) \leq \min_{a \in BR(\bar{\Theta}(s, \pi), s)} u_1(\theta, s, a) \leq u_1(\theta, \pi)$. ■

B. Co-DI

PROOF OF PROPOSITION 3:

Fix $s \in S$. We will argue that $\bar{\Theta}(s, \pi) \subseteq \bar{\Theta}^{D1}(s, \pi)$. This, along with the justified response criterion of JCE and the fact that every JCE is a PBE-H, implies that π is co-D1.

If $\Theta^\dagger(s, \pi) \neq \emptyset$, then $\bar{\Theta}(s, \pi) = \Theta^\dagger(s, \pi)$. Let θ be a type such that $\theta \notin \bar{\Theta}^{D1}(s, \pi)$. Then there is some type $\theta' \neq \theta$ such that $D_\theta(s, \pi) \cup D_\theta^0(s, \pi) \subseteq D_{\theta'}(s, \pi)$. This implies that $\theta' \notin \bar{\Theta}(s, \pi)$, so $\bar{\Theta}(s, \pi) \subseteq \bar{\Theta}^{D1}(s, \pi)$ follows. If $\Theta^\dagger(s, \pi) = \emptyset$, by Lemma A1, $u_1(\theta, s, a) < u_1(\theta, \pi)$ for all $a \in BR(\Theta, s)$. Thus $\bar{\Theta}^{\dagger, D1}(s, \pi) = \emptyset$ as $D_\theta(s, \pi) \cup D_\theta^0(s, \pi) \subseteq D_{\theta'}(s, \pi)$ for all $\theta, \theta' \in \Theta$. Thus, $\bar{\Theta}(s, \pi) = \Theta = \bar{\Theta}^{D1}(s, \pi)$. ■

C. NWBR

LEMMA A2: $\Theta^\ddagger(s, \pi) \subseteq \Theta^\dagger(s, \pi)$ for all $s \in S$ and $\pi \in \Pi$.

PROOF:

If $\theta \notin \Theta^\dagger(s, \pi)$, then by definition $\tilde{D}_\theta^0(s, \pi) \subseteq \bigcup_{\theta' \neq \theta} \tilde{D}_{\theta'}(s, \pi)$. For $\alpha \in D_\theta^0(s, \pi)$, $\alpha \in MBR(\Theta, s) \subseteq \Delta(BR(\Theta, s))$ and $u_1(\theta, s, \alpha) = u_1(\theta, \pi)$, so $\alpha \in \tilde{D}_\theta^0(s, \pi)$. Since $\tilde{D}_\theta^0(s, \pi) \subseteq \bigcup_{\theta' \neq \theta} \tilde{D}_{\theta'}(s, \pi)$, there is some $\theta' \neq \theta$ such that $u_1(\theta', s, \alpha) > u_1(\theta', \pi)$, or equivalently $\alpha \in D_{\theta'}(s, \pi)$. As α is an arbitrary element of $D_\theta^0(s, \pi)$, we conclude that $D_\theta^0(s, \pi) \subseteq \bigcup_{\theta' \neq \theta} D_{\theta'}(s, \pi)$, so $\theta \notin \Theta^\ddagger(s, \pi)$. ■

LEMMA A3: If π is a PBE-H that satisfies NWBR, then, for every $s \in S$, either

- (i) $\Theta^\ddagger(s, \pi) \neq \emptyset$, or
- (ii) $u_1(\theta, s, a) < u_1(\theta, \pi)$ for all $\theta \in \Theta$ and $a \in BR(\Theta, s)$.

The proof of Lemma A3 is analogous to that of Lemma A1, and is given in online Appendix Section OA.5.

PROOF OF PROPOSITION 4:

Let π be a PBE-H that satisfies NWBR, and for every off-path s , let $\alpha_s \in MBR(\hat{\Theta}(s, \pi), s)$ be such that $u_1(\theta, s, \alpha) \leq u_1(\theta, \pi)$ for all $\theta \in \Theta$. We will show that $\hat{\Theta}(s, \pi) \subseteq \bar{\Theta}(s, \pi)$ for all s , so the profile $\tilde{\pi} = (\pi_1, \tilde{\pi}_2)$ in which $\tilde{\pi}_2$ coincides with π_2 for all on-path s and dictates α_s for all off-path s is a JCE that is path equivalent to π .

If $\Theta^\dagger(s, \pi) \neq \emptyset$, then by Lemma A2, $\Theta^\dagger(s, \pi) \subseteq \Theta^\dagger(s, \pi)$, so $\hat{\Theta}(s, \pi) \subseteq \bar{\Theta}(s, \pi)$. If $\Theta^\dagger(s, \pi) = \emptyset$, then by Lemma A3, $u_1(\theta, s, a) < u_1(\theta, \pi)$ for all $\theta \in \Theta$ and $a \in BR(\Theta, s)$, so $\Theta^\dagger(s, \pi) = \emptyset$ and $\hat{\Theta}(s, \pi) = \bar{\Theta}(s, \pi) = \Theta$. ■

PROOF OF LEMMA 4:

Fix PBE-H π . We show that, for all $s \in S$ and $\theta \in \Theta$, $\tilde{D}_\theta(s, \pi) \cup \tilde{D}_\theta^0(s, \pi) \not\subseteq \bigcup_{\theta' \neq \theta} \tilde{D}_{\theta'}(s, \pi)$ if and only if $D_\theta^0(s, \pi) \not\subseteq \bigcup_{\theta' \neq \theta} D_{\theta'}(s, \pi)$. This means that $\Theta^\dagger(s, \pi) = \Theta^\ddagger(s, \pi)$, which implies that $\bar{\Theta}(s, \pi) = \hat{\Theta}(s, \pi)$.

Suppose that $D_\theta^0(s, \pi) \not\subseteq \bigcup_{\theta' \neq \theta} D_{\theta'}(s, \pi)$. Then there is some $\alpha \in MBR(\Theta, s)$ such that $u_1(\theta, s, \alpha) = u_1(\theta, \pi)$ and $u_1(\theta', s, \alpha) \leq u_1(\theta', \pi)$ for all $\theta' \neq \theta$. Since $\alpha \in \Delta(BR(\Theta, s))$, this immediately implies that $\tilde{D}_\theta(s, \pi) \cup \tilde{D}_\theta^0(s, \pi) \not\subseteq \bigcup_{\theta' \neq \theta} \tilde{D}_{\theta'}(s, \pi)$.

Suppose that $\tilde{D}_\theta(s, \pi) \cup \tilde{D}_\theta^0(s, \pi) \not\subseteq \bigcup_{\theta' \neq \theta} \tilde{D}_{\theta'}(s, \pi)$. Then there is some $\alpha \in \Delta(BR(\Theta, s))$ such that $u_1(\theta, s, \alpha) \geq u_1(\theta, \pi)$ and $u_1(\theta', s, \alpha) \leq u_1(\theta', \pi)$ for all $\theta' \neq \theta$. Moreover, since π is a PBE-H, there is some $\alpha' \in \Delta(BR(\Theta, s))$ such that $u_1(\theta, s, \alpha') \leq u_1(\theta, \pi)$. By continuity, there exists some $\alpha'' \in MBR(\Theta, s)$ such that $u_1(\theta, s, \alpha'') = u_1(\theta, \pi) \leq u_1(\theta, s, \alpha)$. Because the game is co-monotonic, $u_1(\theta', s, \alpha'') \leq u_1(\theta', s, \alpha) \leq u_1(\theta', \pi)$ holds for all $\theta' \neq \theta$. Thus, $D_\theta^0(s, \pi) \not\subseteq \bigcup_{\theta' \neq \theta} D_{\theta'}(s, \pi)$. ■

APPENDIX B. SUPPORTING RESULTS FOR THEOREM 1

We use the following lemma in several proofs and examples. We omit its proof, which closely follows that of Proposition 5 in Fudenberg and He (2018).

LEMMA B1: Given $\gamma_2 \in [0, 1)$, suppose that $\pi_{\gamma_2} = (\pi_{1, \gamma_2}, \pi_{2, \gamma_2}) = \lim_{k \rightarrow \infty} \lim_{l \rightarrow \infty} \pi_{\gamma_2, k, l}$ for some sequence of steady-state profiles $\pi_{\gamma_2, k, l} \in \Pi^*(g, \delta_k, \gamma_{1, k, l}, \gamma_2)$, where $\lim_{k \rightarrow \infty} \delta_k = 1$ and $\lim_{l \rightarrow \infty} \gamma_{1, k, l} = 1$ for all k . Then, for each $\theta \in \Theta$, $\pi_{1, \gamma_2}(\cdot | \theta)$ puts support only on the (s, m) that are optimal for type θ against π_{2, γ_2} .

PROOF OF LEMMA 1:

Let $\{\pi_{j, k, l} \in \Pi^*(g, \delta_{j, k}, \gamma_{1, j, k, l}, \gamma_{2, j})\}_{j, k, l \in \mathbb{N}}$ be a sequence of steady-state profiles such that $\lim_{j \rightarrow \infty} \lim_{k \rightarrow \infty} \lim_{l \rightarrow \infty} \pi_{j, k, l} = \pi$, where $\lim_{j \rightarrow \infty} \gamma_{2, j} = 1$, $\lim_{k \rightarrow \infty} \delta_{j, k} = 1$ for all j , and $\lim_{l \rightarrow \infty} \gamma_{1, j, k, l} = 1$ for all j, k . By Lemma B1, for every $\theta \in \Theta$, $\pi_{1, \gamma_{2, j}}(\cdot | \theta) = \lim_{k \rightarrow \infty} \lim_{l \rightarrow \infty} \pi_{1, j, k, l}(\cdot | \theta)$ puts support only on signal-message pairs that are best replies to $\pi_{2, \gamma_{2, j}} = \lim_{k \rightarrow \infty} \lim_{l \rightarrow \infty} \pi_{2, j, k, l}$. Combining this with the upper hemi-continuity of optimal play implies that $\pi_1(\cdot | \theta) = \lim_{j \rightarrow \infty} \pi_{1, \gamma_{2, j}}(\cdot | \theta)$ puts support only on signal-message pairs that are best replies to $\pi_2 = \lim_{j \rightarrow \infty} \pi_{2, \gamma_{2, j}}$. ■

PROOF OF LEMMA 2:

Let $q(\theta, s, m) = \lambda(\theta)\pi_1(s, m|\theta)$ be the distribution over (θ, s, m) induced by λ and π_1 , let X^{on} be the set of sender signal-message pairs that occur with positive probability under π , and let $p_{(s,m)}(\theta)$ denote the conditional probability of θ given $(s, m) \in X^{on}$. For $\varepsilon > 0$, let $Q_\varepsilon = \{q' \in \Delta(\Theta \times S \times M) : \max_{(\theta,s,m)} |q'(\theta, s, m) - q(\theta, s, m)| \leq \varepsilon\}$. Because best response correspondences are upper hemicontinuous, there is an $\varepsilon > 0$ such that every receiver whose belief $\tilde{g}_2 \in \Delta(\Delta(\Theta \times S \times M))$ puts probability at least $1 - \varepsilon$ on Q_ε will respond to every $(s, m) \in X^{on}$ with some $a \in BR(p_{(s,m)}, s)$.

Given the non-doctrinaire prior g_2 , Theorem 4.2 of Diaconis and Freedman (1990) implies that there is some $T > 0$ such that a receiver who has lived more than T periods assigns posterior probability of at least $1 - \varepsilon$ to probability distributions q' within $\varepsilon/2$ distance of whatever empirical distribution they have observed. Moreover, by the law of large numbers, for any $\eta > 0$ we can take this T to be such that, with probability at least $1 - \eta/2$, a receiver who has lived more than T periods assigns probability of at least $1 - \varepsilon$ to $Q_{\varepsilon/2}$.

Fix sequences $\{\delta_n\}_{n \in \mathbb{N}}$, $\{\gamma_{1,n}\}_{n \in \mathbb{N}}$, and $\{\gamma_{2,n}\}_{n \in \mathbb{N}}$, and let $\pi_n = (\pi_{1,n}, \pi_{2,n}) \in \Pi^*(g, \delta_n, \gamma_{1,n}, \gamma_{2,n})$ be a sequence of steady-state profiles such that $\lim_{n \rightarrow \infty} \gamma_{2,n} = 1$ and $\lim_{n \rightarrow \infty} \pi_{1,n} = \pi_1$. The share of receivers in the population who have lived more than T periods is $\gamma_{2,n}^T$, which converges to 1 as $n \rightarrow \infty$. Moreover, $q_n(\theta, s, m) = \lambda(\theta)\pi_{1,n}(s, m|\theta) \rightarrow q$ as $n \rightarrow \infty$. Thus, for every $(s, m) \in X^{on}$ and $\eta > 0$, there exists some $N \in \mathbb{N}$ such that $\pi_{2,n}(BR(p_{(s,m)}, s)|s, m) \geq 1 - \eta$ for all $n > N$. ■

PROOF OF LEMMA 3:

Let $\{\pi_{j,k,l} \in \Pi^*(g, \delta_{j,k}, \gamma_{1,j,k,l}, \gamma_{2,j})\}_{j,k,l \in \mathbb{N}}$ be a sequence of steady-state profiles such that $\lim_{j \rightarrow \infty} \lim_{k \rightarrow \infty} \lim_{l \rightarrow \infty} \pi_{j,k,l} = \pi$, where $\lim_{j \rightarrow \infty} \gamma_{2,j} = 1$, $\lim_{k \rightarrow \infty} \delta_{j,k} = 1$ for all j , and $\lim_{l \rightarrow \infty} \gamma_{1,j,k,l} = 1$ for all j, k . Since $u_1(\theta, s, \pi_2(\cdot|s, m_{s,\tilde{\theta}})) < u_1(\theta, \pi)$ for all $\theta \notin \tilde{\Theta}$, Lemma B1 implies that there is some $J \in \mathbb{N}$ such that, for all $\theta' \notin \tilde{\Theta}$ and $j > J$, $\lim_{k \rightarrow \infty} \lim_{l \rightarrow \infty} \pi_{j,k,l}(s, m_{s,\tilde{\theta}}|\theta') = 0$. Receivers who have never observed the signal-message pair $(s, \tilde{\theta})$ played by a type outside of $\tilde{\Theta}$ would respond to this pair with an action belonging to $BR(\tilde{\Theta}, s)$. Thus $\lim_{k \rightarrow \infty} \lim_{l \rightarrow \infty} \pi_{2,j,k,l}(BR(\tilde{\Theta}, s)|s, m_{s,\tilde{\theta}}) = 1$ if $\lim_{k \rightarrow \infty} \lim_{l \rightarrow \infty} \pi_{j,k,l}(s, m_{s,\tilde{\theta}}|\theta') = 0$ for all $\theta' \notin \tilde{\Theta}$. Since this holds for all $j > J$, $\pi_2(BR(\tilde{\Theta}, s)|s, m_{s,\tilde{\theta}}) = 1$. ■

APPENDIX C. A SUFFICIENT CONDITION FOR STABILITY

DEFINITION C1: A signaling game is strictly monotonic if, for all $\theta, \theta' \in \Theta, s \in S$, and $\alpha, \alpha' \in MBR(\Theta, s)$,

- (i) $u_1(\theta, s, \alpha) \geq u_1(\theta, s, \alpha')$ if and only if $u_1(\theta', s, \alpha) \geq u_1(\theta', s, \alpha')$, and
- (ii) $u_1(\theta, s, \alpha) = u_1(\theta, s, \alpha')$ implies $\alpha = \alpha'$.

Here the first condition is exactly the monotonicity of Cho and Sobel (1990). The second condition requires that the sender preference is a strict order on $MBR(\Theta, s)$.

For a given strategy profile π , let X^{on} be the set of on-path signal-message pairs, let $p_{(s,m)}(\theta)$ denote the conditional probability of θ given $(s,m) \in X^{on}$, let S^{on} be the set of on-path signals, and let S^{off} be the set of off-path signals.

DEFINITION C2: *The JCE π is uniformly justified if*

- (i) *For all $\theta \in \Theta$, there is some $s_\theta \in S$ such that $\max_{m \in M} u_1(\theta, s_\theta, \pi_2(\cdot | s_\theta, m)) > \max_{s \neq s_\theta, m \in M} u_1(\theta, s, \pi_2(\cdot | s, m))$;*
- (ii) *For every $x = (s, m) \in X^{on}$, there is some $a_x \in A$ such that $u_2(p_{(s,m)}, s, a_x) > \max_{a \neq a_x} u_2(p_{(s,m)}, s, a)$;*
- (iii) *For all $s \in S^{off}$, $u_1(\theta, s, a) < u_1(\theta, \pi)$ for all $\theta \in \Theta$ and $a \in BR(\bar{\Theta}(s, \pi), s)$.*

Condition (i) says that every sender type plays exactly one signal and that they have strict incentives to do so. Condition (ii) says that the receiver has a strictly optimal action in response to every on-path signal-message pair. Condition (iii) says that all types are strictly deterred from playing any off-path signal for any justified response.

PROPOSITION C1: *If π is a uniformly justified JCE in a strictly monotonic signaling game, it induces the same distribution over $\Theta \times S \times A$ as a stable profile for all non-doctrinaire priors g_1, g_2 , including those that do not satisfy initial trust.*

Section OA.4 in the online Appendix contains the proof of Proposition C1. Because π is uniformly justified, there is a receiver behavior strategy that makes each type strictly prefer to play their corresponding signal in π , and, when each type does so, leads to the same distribution over $\Theta \times S \times A$ as π . The proof modifies the aggregate response correspondences so that the receiver response matches this behavior strategy with high probability whenever the aggregate sender play is such that some type gives their corresponding signal in π too little probability. Lemma B1 implies the aggregate sender play given by the fixed points of the modified aggregate response correspondence is optimal in the iterated limit. The modification to the receiver aggregate response thus ensures that the limit aggregate sender strategy uses the signals prescribed by π with high probability.³⁰ Additionally, by strict monotonicity and the optimality of the aggregate sender play, the receiver response to any on-path signal-message pair must only depend on the signal, and because receivers strictly prefer to conform to π , the receiver response to any on-path signal-message pair matches the response in π . We show that this, along with the fact that π is uniformly justified, implies that, in the limit, each sender type uses the same distribution over signals as in π_1 . Consequently, the modified aggregate receiver response matches the true aggregate receiver response, and the fixed points of the modified

³⁰Fudenberg and Levine (2006) and Fudenberg and He (2020) proved that some strategy profiles are stable by considering priors that assign high probability to a neighborhood of the target profile. Modifying the aggregate response mapping lets us prove stability for a broad class of priors.

response mapping are valid steady-state profiles that in the limit induce the same distribution over $\Theta \times S \times A$ as π .

APPENDIX D. DETAILS OMITTED FROM SECTION II

Strategy Mapping.—The map $\sigma^{\delta, \gamma_1} : (\Delta(\mathcal{H}_1))^\Theta \times \Delta(\mathcal{H}_2) \rightarrow \Pi_1 \times \Pi_2$ taking the state in period t to the aggregate strategy profile has component mappings $\sigma_1^{\delta, \gamma_1} : (\Delta(\mathcal{H}_1))^\Theta \rightarrow \Pi_1$ and $\sigma_2 : \Delta(\mathcal{H}_2) \rightarrow \Pi_2$ given by $\sigma_1^{\delta, \gamma_1}(\mu_1)[s, m | \theta] = \sum_{h_1: \mathbf{x}_\theta^{\delta, \gamma_1}(h_1) = (s, m)} \mu_\theta[h_1]$ and $\sigma_2(\mu_2)[a | s, m] = \sum_{h_2: \mathbf{y}(s, m | h_2) = a} \mu_2[h_2]$.

Update Rule.—The rule that maps the state in period t to the state in period $t + 1$, $\mathbf{f}^{\delta, \gamma_1, \gamma_2} : (\Delta(\mathcal{H}_1))^\Theta \times \Delta(\mathcal{H}_2) \rightarrow (\Delta(\mathcal{H}_1))^\Theta \times \Delta(\mathcal{H}_2)$, has the following components: the mapping $\mathbf{f}_\theta^{\delta, \gamma_1} : (\Delta(\mathcal{H}_1))^\Theta \times \Delta(\mathcal{H}_2) \rightarrow \Delta(\mathcal{H}_1)$ is given by $\mathbf{f}_\theta^{\delta, \gamma_1}(\mu)[\emptyset] = 1 - \gamma_1$, and $\mathbf{f}_\theta^{\delta, \gamma_1}(\mu)[(h_1, (s, m, a))] = \gamma_1 \mu_\theta[h_1] i_\theta^{\delta, \gamma_1}(h_1, s, m) \sigma_2(\mu)[a | s, m]$, where $(h_1, (s, m, a)) \in \mathcal{H}_1$ is the concatenation of the history $h_1 \in \mathcal{H}_1$ with a period where the sender plays (s, m) and the receiver responds with a , and $i_\theta^{\delta, \gamma_1}(h_1, s, m)$ equals one if a type θ sender with history h_1 plays (s, m) under policy \mathbf{x}_θ and equals zero otherwise. Likewise, $\mathbf{f}_2^{\delta, \gamma_1, \gamma_2} : (\Delta(\mathcal{H}_1))^\Theta \times \Delta(\mathcal{H}_2) \rightarrow \Delta(\mathcal{H}_2)$ is given by $\mathbf{f}_2^{\delta, \gamma_1, \gamma_2}(\mu)[\emptyset] = 1 - \gamma_2$, and $\mathbf{f}_2^{\delta, \gamma_1, \gamma_2}(\mu)[(h_2, (\theta, s, m))] = \gamma_2 \mu_2[h_2] \lambda(\theta) \sigma_1^{\delta, \gamma_1}(\mu)[s, m | \theta]$, where $(h_2, (\theta, s, m)) \in \mathcal{H}_2$ is the concatenation of the history $h_2 \in \mathcal{H}_2$ with a period where the receiver is matched with a type θ sender who plays (s, m) .

Aggregate Response Mapping.—To define the aggregate response mapping, we first define mappings $\mathcal{L}_1^{\delta, \gamma_1} : \Pi_2 \rightarrow (\Delta(\mathcal{H}_1))^\Theta$ and $\mathcal{L}_2^{\gamma_2} : \Pi_1 \rightarrow \Delta(\mathcal{H}_2)$, which output the resulting $t \rightarrow \infty$ limit of the distribution of histories in the sender and receiver populations when the aggregate play of the opposing population is held fixed at a given behavior strategy. For each $\theta \in \Theta$, $\mathcal{L}_1^{\delta, \gamma_1}(\pi_2)[\emptyset] = 1 - \gamma_1$ is the share of type θ senders with the null history. The share of type θ senders with histories $h_{1,t}$ of length $t > 0$ is defined by induction: for each $h_{1,t-1}$, we pair the signal-message pair a type θ sender with history $h_{1,t-1}$ would use with the aggregate receiver strategy π_2 to compute the distribution of period- t outcomes (s_t, m_t, a_t) these senders observe, and assign the corresponding probabilities to the concatenation of these period- t outcomes and $h_{1,t-1}$. (We do this formally in Section OA.2.1 in the online Appendix.) Likewise, $\mathcal{L}_2^{\gamma_2}(\pi_1)[\emptyset] = 1 - \gamma_2$ is the share of receiver agents with the null history. A similar induction procedure gives the share of receiver agents with various histories of length $t > 0$: for each history of length $t - 1$, we take the strategy these agents would use, pair this with the distribution of sender types λ and the aggregate sender strategy π_1 to compute the distribution of period- t outcomes (θ, s_t, m_t) these agents observe, and assign the corresponding probability to the concatenation of the period- t outcomes and the previous history.

The components of the aggregate response mapping $\mathcal{R}^{\delta, \gamma_1, \gamma_2}(\pi) = (\mathcal{R}_1^{\delta, \gamma_1}(\pi_2), \mathcal{R}_2^{\gamma_2}(\pi_1))$ are then found by composing $\mathcal{L}_1^{\delta, \gamma_1}$ and $\mathcal{L}_2^{\gamma_2}$ with strategy mapping $\sigma^{\delta, \gamma_1}$: the aggregate sender response mapping is given by $\mathcal{R}_1^{\delta, \gamma_1}(\pi_2) = \sigma_1^{\delta, \gamma_1}(\mathcal{L}_1^{\delta, \gamma_1}(\pi_2))$, and the aggregate receiver response mapping is given by $\mathcal{R}_2^{\gamma_2}(\pi_1) = \sigma_2(\mathcal{L}_2^{\gamma_2}(\pi_1))$.

REFERENCES

- Abeler, Johannes, Daniele Nosenzo, and Collin Raymond. 2019. "Preferences for Truth-Telling." *Econometrica* 87 (4): 1115–53.
- Banks, Jeffrey S., and Joel Sobel. 1987. "Equilibrium Selection in Signaling Games." *Econometrica* 55 (3): 647–61.
- Battigalli, Pierpaolo. 1987. "Comportamento Razionale ed equilibrio nei giochi e nelle situazioni sociali." http://didattica.unibocconi.it/mypage/upload/48808_20120403_091802_1987THESIS.PDF.
- Battigalli, Pierpaolo, Simone Cerreia-Vioglio, Fabio Maccheroni, and Massimo Marinacci. 2015. "Self-Confirming Equilibrium and Model Uncertainty." *American Economic Review* 105 (2): 646–77.
- Battigalli, Pierpaolo, Alejandro Francetich, Giacomo Lanzani, and Massimo Marinacci. 2019. "Learning and Self-Confirming Long-Run Biases." *Journal of Economic Theory* 183: 740–85.
- Binmore, Ken, and Larry Samuelson. 1999. "Evolutionary Drift and Equilibrium Selection." *Review of Economic Studies* 66 (2): 363–93.
- Blume, Andreas, and Joel Sobel. 1995. "Communication-Proof Equilibria in Cheap-Talk Games." *Journal of Economic Theory* 65 (2): 359–82.
- Chen, Ying, Navin Kartik, and Joel Sobel. 2008. "Selecting Cheap-Talk Equilibria." *Econometrica* 76 (1): 117–36.
- Cho, In-Koo, and David M. Kreps. 1987. "Signaling Games and Stable Equilibria." *Quarterly Journal of Economics* 102 (2): 179–221.
- Cho, In-Koo, and Joel Sobel. 1990. "Strategic Stability and Uniqueness in Signaling Games." *Journal of Economic Theory* 50 (2): 381–413.
- Clark, Daniel. 2020. "Robust Neologism Proofness in Signaling Games." <http://economics.mit.edu/files/16758>.
- Dekel, Eddie, Drew Fudenberg, and David K. Levine. 1999. "Payoff Information and Self-Confirming Equilibrium." *Journal of Economic Theory* 89 (2): 165–85.
- DeMarzo, Peter, and Darrell Duffie. 1999. "A Liquidity-Based Model of Security Design." *Econometrica* 67 (1): 65–99.
- DeMarzo, Peter M., Ilan Kremer, and Andrzej Skrzypacz. 2005. "Bidding with Securities: Auctions and Security Design." *American Economic Review* 95 (4): 936–59.
- Diaconis, Persi, and David Freedman. 1990. "On the Uniform Consistency of Bayes Estimates for Multinomial Probabilities." *The Annals of Statistics* 18 (3): 1317–27.
- Esponda, Ignacio. 2013. "Rationalizable Conjectural Equilibrium: A Framework for Robust Predictions." *Theoretical Economics* 8 (2): 467–501.
- Farrell, Joseph. 1993. "Meaning and Credibility in Cheap-Talk Games." *Games and Economic Behavior* 5 (4): 514–31.
- Fudenberg, Drew, and Kevin He. 2018. "Learning and Type Compatibility in Signaling Games." *Econometrica* 86 (4): 1215–55.
- Fudenberg, Drew, and Kevin He. 2020. "Payoff Information and Learning in Signaling Games." *Games and Economic Behavior* 120: 96–120.
- Fudenberg, Drew, and Kevin He. 2021. "Player-Compatible Learning and Player-Compatible Equilibrium." *Journal of Economic Theory* 194: Article 105238.
- Fudenberg, Drew, and Yuichiro Kamada. 2015. "Rationalizable Partition-Confirmed Equilibrium." *Theoretical Economics* 10 (3): 775–806.
- Fudenberg, Drew, and Yuichiro Kamada. 2018. "Rationalizable Partition-Confirmed Equilibrium with Heterogeneous Beliefs." *Games and Economic Behavior* 109: 364–81.
- Fudenberg, Drew, and David M. Kreps. 1988. "A Theory of Learning, Experimentation, and Equilibrium in Games." <http://economics.mit.edu/files/20982>.
- Fudenberg, Drew, and David M. Kreps. 1994. "Learning in Extensive-Form Games, II: Experimentation and Nash Equilibrium." <http://economics.mit.edu/files/20979>.
- Fudenberg, Drew, and David K. Levine. 1993. "Steady State Learning and Nash Equilibrium." *Econometrica* 61 (3): 547–73.
- Fudenberg, Drew, and David K. Levine. 2006. "Superstition and Rational Learning." *American Economic Review* 96 (3): 630–51.
- Fudenberg, Drew, and Jean Tirole. 1991a. *Game Theory*. Cambridge, MA: MIT Press.
- Fudenberg, Drew, and Jean Tirole. 1991b. "Perfect Bayesian Equilibrium and Sequential Equilibrium." *Journal of Economic Theory* 53 (2): 236–60.

- Gittins, John C.** 1979. "Bandit Processes and Dynamic Allocation Indices." *Journal of the Royal Statistical Society, Series B* 41 (2): 148–64.
- Goltsman, Maria, and Gregory Pavlov.** 2011. "How to Talk to Multiple Audiences." *Games and Economic Behavior* 72 (1): 100–122.
- Gordon, Sidartha, Navin Kartik, Melody Pei-yu Lo, Wojciech Olszewski, and Joel Sobel.** 2021. "Effective Communication in Cheap Talk Games." <http://www.columbia.edu/~nk2339/Papers/GKLOS-EffectiveCommunication.pdf>.
- Hart, Sergiu.** 2002. "Evolutionary Dynamics and Backward Induction." *Games and Economic Behavior* 41 (2): 227–64.
- Jehiel, Philippe, and Dov Samet.** 2005. "Learning to Play Games in Extensive Form by Valuation." *Journal of Economic Theory* 124 (2): 129–48.
- Kalai, Ehud, and Ehud Lehrer.** 1993. "Rational Learning Leads to Nash Equilibrium." *Econometrica* 61 (5): 1019–45.
- Kartik, Navin.** 2009. "Strategic Communication with Lying Costs." *Review of Economic Studies* 76 (4): 1359–95.
- Kartik, Navin, Marco Ottaviani, and Francesco Squintani.** 2007. "Credulity, Lies, and Costly Talk." *Journal of Economic Theory* 134 (1): 93–116.
- Kohlberg, Elon, and Jean-Francois Mertens.** 1986. "On the Strategic Stability of Equilibria." *Econometrica* 54 (5): 1003–37.
- Lehrer, Ehud, and Eilon Solan.** 2007. "Learning to Play Partially-Specified Equilibrium." <http://www.math.tau.ac.il/~lehrer/Papers/LearningPSCE-web.pdf>.
- Lipman, Barton L., and Duane J. Seppi.** 1995. "Robust Inference in Communication Games with Partial Provability." *Journal of Economic Theory* 66 (2): 370–405.
- Matthews, Steven A., Masahiro Okuno-Fujiwara, and Andrew Postlewaite.** 1991. "Refining Cheap-Talk Equilibria." *Journal of Economic Theory* 55 (2): 247–73.
- Nachman, David C., and Thomas H. Noe.** 1994. "Optimal Design of Securities under Asymmetric Information." *Review of Financial Studies* 7 (1): 1–44.
- Nöldeke, Georg, and Larry Samuelson.** 1993. "An Evolutionary Analysis of Backward and Forward Induction." *Games and Economic Behavior* 5 (3): 425–54.
- Olszewski, Wojciech.** 2006. "Rich Language and Refinements of Cheap-Talk Equilibria." *Journal of Economic Theory* 128 (1): 164–86.
- Rabin, Matthew.** 1990. "Communication between Rational Agents." *Journal of Economic Theory* 51 (1): 144–70.
- Rabin, Matthew, and Joel Sobel.** 1996. "Deviations, Dynamics, and Equilibrium Refinements." *Journal of Economic Theory* 68 (1): 1–25.
- Riley, John G.** 1979. "Informational Equilibrium." *Econometrica* 47 (2): 331–59.
- Rubinstein, Ariel, and Asher Wolinsky.** 1994. "Rationalizable Conjectural Equilibrium: Between Nash and Rationalizability." *Games and Economic Behavior* 6 (2): 299–311.
- Sobel, Joel, Lars Stole, and Iñigo Zapater.** 1990. "Fixed-Equilibrium Rationalizability in Signaling Games." *Journal of Economic Theory* 52 (2): 304–31.
- Zapater, Iñigo.** 1997. "Credible Proposals in Communication Games." *Journal of Economic Theory* 72 (1): 173–97.