Noname manuscript No.
(will be inserted by the editor)

On Defining Ex Ante Payoffs in Games with Diffuse Prior

Attila Ambrus · Aaron Kolb

Received: date / Accepted: date

Abstract While the diffuse prior has been widely used in applied economic theory for its technical convenience and as a way of modeling complete lack of knowledge, it is not formally defined, nor are ex ante payoffs in games under this prior. In this paper, we provide a formal treatment of the diffuse prior which can validate its application in games. We consider stationary games, in which players' signals are translation invariant in the true state and players' payoffs are translation invariant in actions together with the state. We show that strategies which admit well-defined expected payoffs under the diffuse prior are essentially stationary, being almost translation invariant in signals. Our analysis builds on two formal definitions. We define the diffuse prior through a limit construction, using sequences of well-defined priors that become increasingly dispersed. A class of strategy profiles is admissible if for any strategy profile, each player's ex ante payoff along these sequences converges to a limit that does not depend on the particular sequence. A secondary contribution of the paper is an extension of the concept of distributional strategies (Milgrom and Weber 1985) to a class of multistage games.

Keywords diffuse prior · stationary games · distributional strategies

Acknowledgements The authors thank Volodymyr Baranovskyi, Joel Sobel, Tom Wiseman, and anonymous referees for helpful comments and encouragement. The usual disclaimer applies.

JEL Classification C72, D80

1 Introduction

The diffuse, or uninformative, prior is often interpreted informally as a uniform distribution on the real line. This prior has two advantages for use in economics: first, as it represents complete ignorance, it is appropriate for modeling situations in which agents have no advance knowledge of the environment; and second, it makes updating beliefs through Bayes' rule computationally simpler. This tractability comes from ex ante symmetry across all states. When signals are the sum of the state and a conditionally i.i.d. noise, posterior beliefs about the state given two signal values are simply translations of one another. However, although the diffuse prior is commonly used, it is not formally defined: any uniform distribution must have constant density and must integrate to one, but any positive constant density, integrated over the real line, yields infinity, and zero density integrates to zero. This lack of a formal representation means that ex ante expected payoffs are not defined when driven by a random variable drawn from a diffuse prior distribution. The existing literature has circumvented this issue by leaving expected payoffs undefined and instead focusing on payoffs conditional on signal realizations (for example, Friedman (1991), Klemperer (1999), Morris and Shin (2002, 2003), and Myatt and Wallace (2014)).

In this paper, we develop a method for formally defining expected payoffs under a diffuse prior, and thereby bringing them into the realm of traditional game theory, where expected payoffs are assumed to

Attila Ambrus
Duke University
E-mail: attila.ambrus@duke.edu

Aaron Kolb
Indiana University Kelley School of Business
E-mail: kolba@indiana.edu

be well-defined for all strategy profiles of a game. This can be directly useful in applications for various reasons. One might be interested in ex ante (or behind the veil of ignorance) welfare evaluation. Or in some contexts it is natural to allow for an ex ante participation constraint for one of the players. For example, when a management consultant is hired to examine and improve a firm's performance, the management consultant has to decide whether to accept the task before inspecting the firm's internal structure and production. In other contexts an agent might need her strategy approved by a principal ex ante, before acquiring private information, which only leads to a well-defined problem if the principal can form ex ante payoff expectations. Our method validates the conventional but informal approach of handling a diffuse prior, and providing rigorous foundations for diffuse prior might encourage more applied theory research to feature analytically tractable games with diffuse prior.

We claim that the diffuse prior can be rigorously constructed as a limit of well-defined distributions, and that expected payoffs under a diffuse prior can be defined in certain cases. We define a class of games with stationary information and payoff structures, where signals and payoffs are translation invariant in the following way: the distribution of signals is translation invariant in the true state, and payoffs are translation invariant in all actions together with the true state. We then show that as long as a class of strategy profiles includes all stationary strategies, then admissibility requires that all strategies are nearly stationary in a precise sense.

We capture the main features of this uninformative prior by using a sequence of (proper) measures that *diffuse* in a formal sense (Definition 4). We say that a class of strategies is *admissible* (Definition 5) when ex ante expected payoffs, taken along any diffusing sequence of proper priors, have a well-defined limit that does not depend on the particular sequence.[1] Stationary strategies are unsurprisingly admissible, as given any stationary strategy profile, expected payoffs conditional on all signal realizations are the same. Our main result (Theorem 1) states roughly that in any class of admissible strategy profiles that includes all stationary strategies, every strategy is *nearly stationary* in a particular sense (Definition 3). Furthermore, every such strategy is payoff-equivalent to some stationary strategy. We extend this result to a class of multistage games in Section 4.

An interpretation of our results is as follows. We offer a limit construction of the diffuse prior which allows ex ante payoffs under this prior to be defined as limits of payoffs under proper priors. However, the existence of such ex ante payoffs places limits on the strategy profiles available. To the extent that formal equilibrium concepts or welfare analyses require well-defined ex ante payoffs, it is useful to identify an admissible class of strategy profiles within which ex ante payoffs are guaranteed to be well-defined, and hence we define admissibility as a property of strategy profiles. In stationary games it is natural to assume that strategy sets include stationary strategies, and hence we require an admissible class to include all stationary strategies; in other words, strategies must yield well-defined payoffs when played against stationary strategies. This notion of admissibility then implies that all strategies are close to stationary strategies.[2]

We demonstrate the applicability of our results in two contexts: (i) in the context of beauty contest games introduced in Morris and Shin (2002),[3] and (ii) in the delegation framework of Ambrus et al. (2019).

A secondary contribution of our paper is an extension of the concept of distributional strategies (Milgrom and Weber 1985) to a class of multistage games. The key additional feature of distributional strategies for multistage games is that a player's actions in later stages can depend on actions by other players in earlier stages, and thus the distributional strategy includes dimensions for past actions.

Before proceeding, we briefly comment on the existing literature in statistics and probability theory on the subject of non-informative priors. Analysts have long debated the best way to impose a prior belief when performing parameter estimation and have recognized various desirable features. Laplace (1951) argued for a uniform or flat prior, from a principle of insufficient reason: without further information, any two possible values of the parameter should be considered equally likely. Jeffreys (1946, 1961) proposed a particular rule for selecting a prior, later known as the *Jeffreys prior*, for a given data-generating process

[1] In a paper largely unrelated to our work, Dale and Morgan (2015) consider specific sequences of proper priors diffusing in a similar sense as in our definition, in the context of a specific game from Morris and Shin (2002). They do not investigate the possibility of defining ex ante expected payoffs in the game; instead they are interested in comparing equilibrium predictions of the model with proper versus improper priors. Equilibrium (and in general, strategic) analysis is not part of the current paper.

[2] In Section 6, we explore alternative notions of admissibility.

[3] See also Hellwig and Veldkamp (2009) and Angeletos et al. (2010).

which would be invariant to reparameterizations.[4] While the Jeffreys prior need not be uniform, in the case of data drawn from a unidimensional normal distribution with known variance and unknown mean (as in many economic applications, including those considered in this paper), the Jeffreys prior is indeed uniform. (Of course, over the real line, this results in an improper prior.) Hence our use of a diffuse prior is consistent both with economic applications and with both approaches mentioned above. We refer the reader to Kass and Wasserman (1996) for an extensive review of the literature on non-informative priors, and to Yang and Berger (1998) for a catalog of such priors.

The paper is structured as follows. Section 2 gives an overview of the main ideas and outlines the steps of the main proof. Section 3 gives the formal analysis for single-stage games, and Section 4 extends this analysis to multistage games. Section 5 shows how our results can be applied to a beauty contest game and to a companion paper on delegation. Section 6 discusses alternative approaches to handling nonstationary strategies, and Section 7 concludes. The appendix contains proofs not provided in the body of the paper.

2 Overview

To capture the diffuse prior as a limit object, we define sequences of proper measures to be diffusing if, roughly, the measures become increasingly uniform and spread out over the real line. Our definition allows for a large class of diffusing sequences, including sequences of uniform distributions on $[-n, n]$ or sequences of normal distributions with variance n, with n going to infinity.

We will define payoffs under a diffuse prior in cases where the limit of payoffs taken along any diffusing sequence exists and is independent of the sequence. Before defining the class of games we consider, we demonstrate how some concrete functions from $\mathbb{R}$ to $\mathbb{R}$ stand up to this criterion. Clearly, a constant function, when integrated with respect to any probability measure, integrates to that constant, and so all diffusing sequences result in the same limit, and thus a constant function is admissible. In addition, the function $x \mapsto \mathbb{1}_{[0,1]}(x)$, which takes the value 1 if $x \in [0, 1]$ and 0 otherwise, is also admissible; it is not difficult to show that along any diffusing sequence, the expected value approaches 0. On the other hand, a function like $x \mapsto \mathbb{1}_{[0,\infty)}$ is *not* admissible. One could obtain a limit of 0 by defining a diffusing sequence $(\mathbb{P}^1_n)_{n\in\mathbb{N}}$ with the densities $\frac{\mathbb{1}_{[-n^2,n]}}{n^2+n}$ and a limit of 1 by a different diffusing sequence $(\mathbb{P}^2_n)_{n\in\mathbb{N}}$ with densities $\frac{\mathbb{1}_{[-n,n^2]}}{n^2+n}$. The key property of admissible functions here is that they are constant or "nearly" constant in some formal sense. As we are interested in games of asymmetric information, and not real-valued functions per se, the exercise is more subtle than the above examples suggest. Nonetheless, the above intuition plays a key role in the analysis that follows.

In the baseline model, we analyze static n-player games with asymmetric information. At the beginning of the game, the state of the world $\theta \in \mathbb{R}$ is drawn according to a diffuse prior, and players receive private, conditionally i.i.d. signals s_i about θ. Players then simultaneously choose real-valued actions. We assume that the game has a stationary structure in terms of signals and payoffs: (i) for each player there is some distribution F_i such that for all θ, $s_i - \theta$ is drawn from this distribution and (ii) payoffs are invariant to a translation of all actions and θ by a constant. The model also accommodates uninformed players.

Strategies must specify (distributions over) actions given a player's private signal. Since we start from strategy sets that are not restricted to be stationary, we need to provide a flexible and careful formal definition of strategies. Since the state space is uncountable, it is not practical to define strategies as products of signal-dependent distributions over actions. The key tension is that desirable topologies should be both rich enough so that payoff functions are continuous in strategies, but also coarse enough so that the strategy space is compact. We follow the approach of Milgrom and Weber (1985) in using *distributional strategies*, which are measures μ over the product space of signals and actions. A distributional strategy induces a conditional distribution on actions for any given signal, and by integrating over signals, it induces a conditional distribution on actions for any given θ.

It is convenient to normalize distributional strategies by setting the marginal distribution over the signal dimension to be some fixed, arbitrary distribution G with full support on the real line.[5] To describe behavior and payoffs conditional on arbitrary θ, it is useful to exploit the stationary structure of the

[4] For a parameter θ and data $X \sim f(x|\theta)$, the prior density is, up to a scaling factor, $\det(\mathbf{I}(\theta))^{1/2}$, where $I(\theta)$ is the Fisher information matrix $\mathbf{I}(\theta)_{ij} = \mathbb{E}\left[-\frac{\partial^2 l}{\partial\theta_i \partial\theta_j}\right]$ and where l is the log-likelihood ratio.

[5] Here we emphasize that G is not necessarily the actual distribution over player types, in a departure from Milgrom and Weber (1985). See the discussion following Definition 1.

game and define “recentered” strategies μ_i^x, which, roughly speaking, specify behavior relative to the private signal as if that signal had been shifted down by the constant x. A stationary strategy then, as we define it, is one which is translation invariant, in the sense that if the signal is shifted by a constant, then actions are shifted by the same constant; it follows that for stationary strategies, all of its recentered strategies are the same strategy. Under a mild assumption on the action sets, we show that the space of all distributional strategies, denoted $\mathcal{M}_i$, is compact.[6] We say that a strategy μ_i has a *limit strategy* $\mu_i^* \in \mathcal{M}_i$ if $\lim_{\theta\to-\infty}\mu_i^\theta = \lim_{\theta\to+\infty}\mu_i^\theta = \mu_i^*$ (Definition 3). We show that whenever this property holds, the strategy μ_i^* must be stationary, and hence we call μ *nearly stationary.*

We assume that the game is irreducible in the sense that there are no redundant strategies – there are no two strategies which always yield the same payoffs. Without this assumption, there would be no hope of disciplining the set of admissible strategies, since a player could combine such redundant strategies in arbitrary ways without affecting payoffs.

Our main result, Theorem 1, has two components. First, it says that the class of nearly stationary strategies is admissible, so near stationarity is sufficient for admissibility. Second, it says that in any admissible class of strategies which is at least as large as the set of stationary strategies, all strategies are nearly stationary. Since these strategies are payoff-equivalent to stationary strategies, we argue that for a game with diffuse prior to have well-defined ex ante expected payoffs, essentially all strategies have to be stationary.

We begin the proof of the necessity part of Theorem 1 by establishing the existence of some distribution $\mu_i^* \in \mathcal{M}_i$ with the following property: for all $\eta > 0$, μ_i^θ is within η of μ_i^* for an infinite measure set of θ (Lemma 7). We then call μ_i^* an *attraction.* We show that this is a weaker condition than near stationarity; a necessary condition for μ_i to be near μ_i^* is that μ_i^* is an attraction.

Next, we argue that there can be at most one such attraction for any strategy which is part of an admissible class. The proof of this claim is by contradiction and contains several steps. We suppose that μ_i^* and $\hat{\mu}_i$ are two distinct attractions for player i's strategy μ_i. We argue that by the irreducibility assumption, there must exist some profile of stationary strategies of the rivals, μ_{-i}, against which these distributions yield distinct expected payoffs. Given $\eta > 0$, we can construct a sequence of measures $(\mathbb{P}_n^1)_{n\in\mathbb{N}}$ (resp. $(\mathbb{P}_n^2)_{n\in\mathbb{N}}$) that places increasing mass on θ such that μ_i^θ is within η of μ_i^* (resp. $\hat{\mu}_i$). By continuity and translation, the limit of expected payoffs can be close to either of two distinct values. This violates admissibility, giving the desired contradiction.

Given the unique attraction μ_i^*, we show that μ_i^* is a limit strategy for μ_i. If μ_i has any limit strategy, that strategy must be an attraction, so μ_i^* is the only candidate. We show that if μ_i^* is not a limit strategy for μ_i, then there is a compact set of strategies that does not contain μ_i^* but contains μ_i^θ for a measurable set of infinitely many θ. This compact set itself contains an attraction, and this contradicts the uniqueness of the attraction μ_i^*.

To complete the proof, we argue that μ_i^* is stationary, so that μ_i is nearly stationary.

3 Model

In this section, for expositional reasons, we consider single-stage games; we later extend the results to multistage games in Section 4. Before analyzing the diffuse prior, we specify the class of games we consider, which consists of an information structure and a payoff structure.[7]

3.1 Setup

Let Γ denote the game. There are n players indexed by $i \in \mathcal{I} := \{1, 2, \ldots, N\}$. All players assign a diffuse prior (formalized in Section 3.3) to the state of the world, $\theta \in \Theta := \mathbb{R}$. Players are categorized as either *informed* or *uninformed.* Informed players receive signals $s_i \in S_i := \mathbb{R}$ which are conditionally independent given θ with distributions $s_i - \theta =: \hat{s}_i \sim F_i$ for some cumulative distribution function F_i on $\mathbb{R}$ admitting a positive density f_i; uninformed players do not receive such a signal. We use $\mathcal{I}^{inf}$ and $\mathcal{I}^{un}$ to denote the sets of informed and uninformed players, respectively. Let $S := \times_{i\in\mathcal{I}^{inf}} S_i$ (and likewise for s and $\hat{s}$) and let F be the joint CDF, $F(\hat{s}) = \times_{i\in\mathcal{I}^{inf}} F_i(\hat{s}_i)$. We let X_i denote a copy of $\mathbb{R}$, one for each player; each player simultaneously chooses an action a_i from a set of available actions,

[6] We use the topology of weak convergence for measures and distributions, which is metrized by the Prokhorov distance, d_P, and we use the usual topology for $\mathbb{R}$; continuity and compactness are with respect to these topologies.

[7] Since we are concerned with admissibility, equilibrium plays no role in our analysis, and we do not specify an equilibrium concept.

denoted $A_i(s_i) \subset X_i$ for $i \in \mathcal{I}^{inf}$ and $A_i \subset X_i$ for $i \in \mathcal{I}^{un}$. Given θ and a vector of actions $a = (a_i, a_{-i})$ where $a_{-i} := (a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_N)$, players receive payoffs $u_i(a_i, a_{-i}, \theta)$. We assume that the payoff functions u_i are continuous in all arguments.

Informally, a mixed strategy is a randomization over pure strategies. However, in games of incomplete information with uncountable type spaces (such as in our class of games), topological problems arise from interpreting strategies as mixed strategies (i.e., distributions over pure strategies) or behavioral strategies (i.e., products of history-contingent distributions over actions). Although the standard way to define mixed strategies in games with finite or countable signal spaces is as products, over signals, of signal-dependent distributions, this approach runs into problems when signal spaces are uncountable. As an alternative, Milgrom and Weber (1985) (hereafter "MW") introduce the notion of *distributional strategies*, whereby strategies are defined as joint distributions over signals and actions. We adapt this concept to our needs in Definition 1.

Toward Definition 1, let G be an arbitrary distribution over $\mathbb{R}$ (the common signal space for informed players) with full support, and let ϕ denote the measure induced by this distribution G. Let λ denote the Lebesgue measure. We impose one additional requirement on ϕ, that ϕ and λ are mutually absolutely continuous, that is, they have the same zero-measure sets.[8] We use $\mu_i(\cdot|s_i)$ to denote (versions of) the conditional distribution over actions given signals. We are now ready to define strategies, which for informed players take the form of distributional strategies.

Definition 1 (Strategies) A strategy for a player $i \in \mathcal{I}^{inf}$ is a probability measure μ_i on $S_i \times X_i$ such that:

- (Marginal Distribution over Signals) For all measurable $T \subseteq S_i$, $\mu_i(T \times X_i) = \phi(T)$.
- (Proper Support) For any version of the conditional distribution $\mu_i(\cdot|s_i)$, for $s_i \in S_i$ except on a set of measure zero, the support of $\mu_i(\cdot|s_i)$ is a subset of $A_i(s_i)$.

A strategy for a player $i \in \mathcal{I}^{un}$ is a probability measure μ_i on X_i whose support is a subset of A_i.

A few comments on Definition 1 are in order. The first property is a normalization and represents a departure from MW. The key feature is that the distribution G (which defines ϕ) has full support over the reals. The strategic content of a distributional strategy is in the *conditional* distributions $\mu_i(\cdot|s_i)$ which describe how the player behaves given a particular signal s_i; this content is independent of the particular marginal distribution G. In contrast, MW use the actual distribution over types (here, signals) to play the role of G. We argue that this is not necessary, since the strategic content is contained in the conditional distributions, and the marginal distribution is merely a tool for "packaging" these into a joint distribution. In addition, it is not possible to use the actual type distribution in our setting since, prior to the θ realization, signals for informed players have ex ante diffuse distributions over $\mathbb{R}$. The same applies to our multistage version in Section 4.

The second property in Definition 1 ensures that with probability one, actions are chosen from $A_i(s_i)$. For uninformed players, we can dispense with the signal dimension and define strategies as measures over actions alone.

Under the above definition, a pure strategy for an informed player i is one such that for all s_i, $\mu_i(\cdot|s_i)$ places all mass on a single point (which can depend on s_i), where $\mu_i(\cdot|s_i)$ here denotes the regular conditional probability. That is, for each s_i, there exists $x \in X_i$ such that for all measurable $Y \subseteq X_i$, $\mu_i(Y|s_i) = 1$ if $x \in Y$ and 0 otherwise. (A pure strategy for an uninformed player has the analogous property.)

Since X_i and $S_i \times X_i$ are complete and separable, the spaces of probability measures $\Delta(X_i)$ and $\Delta(S_i \times X_i)$, under the topology of weak convergence, are metrized by the Prokhorov distance.[9] We let $\mathcal{M}_i$ denote the space of all strategies for player i, where $\mathcal{M}_i \subset \Delta(X_i)$ if player i is uninformed and $\mathcal{M}_i \subset \Delta(S_i \times X_i)$ if player i is informed; on each $\mathcal{M}_i$, we use the same Prokhorov distance.

3.2 Stationarity

As the diffuse prior implies symmetry across states, we focus on games where this symmetry holds for all components of the game. We label these games stationary. Below we define stationarity of signals, payoffs and strategies.

[8] For example, if G is the CDF of the standard normal distribution, then G and its associated ϕ satisfy these properties.
[9] See Billingsley (2009, Theorem 6.8).

We assume that the signal structure of the game is stationary so that conditional on any realization θ, the values s_i are drawn i.i.d. with $s_i - \theta \sim F_i$ for some distribution F_i on $\mathbb{R}$ as described earlier. In addition, we assume that payoffs are stationary, in the sense that for any fixed strategy profile of the uninformed players, all players' payoffs are invariant to a shift in the informed players' actions and the state shifted by the same constant: for all $i \in \mathcal{I}$, and strategy profiles a, $u_i(a, \theta) = u_i(a - \theta^{inf}, 0)$, where we use θ^{inf} to denote the vector of length N whose i^{th} component takes the value θ if $i \in \mathcal{I}^{inf}$ and 0 otherwise.[10]

Compactness of the strategy space is useful throughout our analysis. We therefore assume that the action space available after each signal is a compact set; since we are focusing on stationary games, we also assume that action sets after different signals are translations of one another.

Assumption 1 *For any player $i \in \mathcal{I}^{inf}$, the action space available after signal 0, $A_i(0)$, is compact and for all $s_i \neq 0$, $A_i(s_i) = A_i(0) + s_i$. For $i \in \mathcal{I}^{un}$, the action space available A_i is compact.*

For example, suppose action sets are compact intervals, centered at the signal for informed players; specifically, suppose that for each informed player i there exists some constant $M_i > 0$ such that for all s_i, $A_i(s_i) = [s_i - M_i, s_i + M_i]$, and for each uninformed player i there exists $M_i > 0$ such that $A_i = [-M_i, M_i]$. Then Assumption 1 is satisfied. An implication of Assumption 1 with our notion of strategies is the following.

Lemma 1 *Under Assumption 1, each $\mathcal{M}_i$ is compact.*

For informed players, a strategy as defined above gives a distribution over signals and actions with an arbitrary fixed marginal distribution over the signal dimension G. To describe joint distributions over actions and signals (and ultimately, to describe payoffs) conditional on arbitrary θ values, it is useful to define a "recentering function." Formally, for any $i \in \mathcal{I}^{inf}$, strategy $\mu_i \in \mathcal{M}_i$, and $x \in \mathbb{R}$, we define a strategy μ_i^x by specifying its conditional distributions $\mu_i^x(Y|s_i) = \mu_i(Y - x|s_i - x)$ for all $Y \subseteq X_i$ and all $s_i \in S_i$ except on a set of measure zero under ϕ, where $Y - x$ denotes the set Y shifted by $-x$. We then say that μ_i^x is a *recentered strategy*; note that in the case of $x = 0$, we have $\mu_i^0 = \mu_i$. For intuition, it is helpful for a moment to interpret strategies as inducing distributions over *markups*, that is, actions minus signals. Then the strategy μ_i^x gives the same distribution over markups after signal s_i as strategy μ_i does after signal $s_i - x$. Example 1 below gives more detail. For completeness, we define $\mu_i^x = \mu_i$ for uninformed players in this section.

The recentering function also allows us to readily define stationarity of strategies.

Definition 2 (Stationary Strategies) For $i \in \mathcal{I}^{inf}$, a strategy μ_i is stationary if the distribution over i's action and i's private signal are together translation invariant: for all $x \in \mathbb{R}$, $\mu_i = \mu_i^x$. For $i \in \mathcal{I}^{un}$, all strategies are stationary.

A stationary strategy in the case of single-stage games can be described by a single distribution over $A_i(0)$, interpreted as a (possibly random) *markup* x so that the action chosen after receiving signal s_i is $a_i = s_i + x$. As we prove later (see Lemma 5 in the appendix), for an arbitrary strategy μ_i, the map $\theta \mapsto \mu_i^\theta$ from $\mathbb{R}$ to $\mathcal{M}_i$ is uniformly continuous.

Example 1 Suppose the (nonstationary) strategy μ_i specifies

- if $s_i < 1$, play $a_i \sim U[s_i, s_i + 1]$ (i.e., a markup distributed $U[0, 1]$,)
- if $s_i \geq 1$, play $a_i \sim U[s_i + 2, s_i + 3]$ (i.e., a markup distributed $U[2, 3]$.)

The recentered strategy $\mu_i^{1/2}$ (where $x = 1/2$) prescribes the following behavior. If $s_i < 3/2$ is observed, apply a random markup distributed $U[0, 1]$ (as the original strategy μ_i specifies for a signal $s_i - x = s_i - 1/2 < 1$), that is, play $a_i \sim U[s_i, s_i + 1]$; if $s_i \geq 3/2$ is observed, apply a markup distributed $U[2, 3]$ (as μ_i specifies for a signal $s_i - x = s_i - 1/2 \geq 1$), that is, play $a_i \sim U[s_i + 2, s_i + 3]$. Note that μ_i here is not stationary, but given the first bullet point, it would be stationary if for all s_i, $a_i \sim U[s_i, s_i + 1]$.

As mentioned, the recentering function allows us to easily express ex interim expected payoffs (that is, conditional on θ) for arbitrary θ realizations. Since we consider a stationary payoff and signal structure, we can calculate expected payoffs under strategies μ_i conditional on $\theta = x$ as the expected payoffs under strategies μ_i^{-x} conditional on $\theta = 0$ (see equation (1) below). Let μ^x denote the strategy profile where each player $i \in \mathcal{I}$ plays μ_i^x (and in particular, each player $i \in \mathcal{I}^{un}$ plays μ_i).

[10] Likewise, we use θ_{-i}^{inf} to denote the vector of length $N - 1$ formed from θ^{inf} by excluding player i.

A profile of strategies induces ex interim expected payoffs for each player as follows:

$$\begin{aligned}
u_i(\mu,\theta) &:= \int_{\hat{s}\in\mathbb{R}^{|\mathcal{I}^{inf}|}} \int_{a\in X} u_i(a,\theta)\, d((\times_{j\in\mathcal{I}^{inf}}\mu_{j,X_j}(a_j|\hat{s}_j+\theta)) \times (\times_{j\in\mathcal{I}^{un}}\mu_j(a_j)))dF(\hat{s}) \\
&= \int_{\hat{s}\in\mathbb{R}^{|\mathcal{I}^{inf}|}} \int_{a\in X} u_i(a-\theta^{inf},0)\, d((\times_{j\in\mathcal{I}^{inf}}\mu_{j,X_j}(a_j|\hat{s}_j+\theta)) \times (\times_{j\in\mathcal{I}^{un}}\mu_j(a_j)))dF(\hat{s}) \\
&= \int_{\hat{s}\in\mathbb{R}^{|\mathcal{I}^{inf}|}} \int_{a\in X} u_i(a,0)\, d((\times_{j\in\mathcal{I}^{inf}}\mu_{j,X_j}(a_j+\theta|\hat{s}_j+\theta)) \times (\times_{j\in\mathcal{I}^{un}}\mu_j(a_j)))dF(\hat{s}) \\
&= \int_{\hat{s}\in\mathbb{R}^{|\mathcal{I}^{inf}|}} \int_{a\in X} u_i(a,0)\, d((\times_{j\in\mathcal{I}^{inf}}\mu^{-\theta}_{j,X_j}(a_j|\hat{s}_j)) \times (\times_{j\in\mathcal{I}^{un}}\mu_j(a_j)))dF(\hat{s}) \\
&= u_i(\mu^{-\theta},0), \qquad (1)
\end{aligned}$$

where the second line uses stationarity of payoffs, the third line performs a change of variables for each a_j, $j\in\mathcal{I}^{inf}$, the fourth line uses the definition of the recentering function, and the fifth line uses the definition from the first line. Note that if all players play *stationary* strategies, then it can be seen from (1) that they obtain the same interim payoffs for all θ realizations, as $\mu^{-\theta}=\mu$ and hence $u_i(\mu^{-\theta},0)=u_i(\mu,0)$ in that case.

In general, ex post payoffs can vary widely due to the noise in the signals. To ensure that the ex interim payoffs (given a realization of θ, but not the signal realization) are finite, we impose the following assumption, which is a joint condition on the payoff function, action sets and signal distribution. Here, for each noise realization $\hat{s}$, $A(\hat{s}+\theta)$ denotes the set of action profiles a such that $a_j\in A_j(\hat{s}_j+\theta)$ for all $j\in\mathcal{I}^{inf}$ and $a_j\in A_j$ for all $j\in\mathcal{I}^{un}$.

Assumption 2 (Bounded Interim Payoffs) *Together, the signal distributions, action sets, and payoff functions are such that* $\int_{\hat{s}\in\mathbb{R}^{|\mathcal{I}^{inf}|}}\sup_{a\in A(\hat{s}+\theta)}|u_i(a_i,a_{-i},\theta)|dF(\hat{s})$ *is bounded over all* $i\in\mathcal{I}$ *and* $\theta\in\mathbb{R}$.

Assumption 2 holds, for example, when payoff functions are finite polynomial, the signal distributions conditional on the state are normal (which includes the commonly analyzed case of normally distributed signals and quadratic loss functions), and action sets are intervals of bounded size containing the signal. This assumption ensures that expected interim payoffs are bounded.

Lemmas 5 and 6 in the appendix establish useful continuity properties of the recentering function and interim payoffs.

Next we define a weakening of stationarity that we will later prove to be both necessary and sufficient for admissibility. Recall that our arbitrary specification G for the marginal distribution over signals allows us operate with distributional strategies, which are joint distributions over signals and actions. Hence, to describe the distance between strategies, we must use a suitable metric for distances between distributions, and thus we use the Prokhorov metric. Convergence of strategies in the Prokhorov metric means that the measure these strategies assign to well-behaved sets (of pairs of signals and actions) also converge. Toward defining near stationarity, we begin by defining a limit strategy (if it exists) as a strategy that a given strategy approaches in the limit when recentering at extreme states, with distance between strategies being measured by the Prokhorov metric. In other words, the original strategy behaves increasingly similar to the limit strategy when the state becomes very high or very low.

Definition 3 (Limit Strategies and Near Stationarity) For any $i\in\mathcal{I}$ and any strategy μ_i, we say that a strategy μ_i^* is a limit strategy for μ_i if $\lim_{\theta\to-\infty}\mu_i^\theta=\lim_{\theta\to+\infty}\mu_i^\theta=\mu_i^*$, with limits taken with respect to the Prokhorov metric. If μ_i^* is stationary, then we classify μ_i as *nearly stationary.*

Note that stationarity implies near stationarity, as a stationary strategy is its own limit strategy. In particular, for $i\in\mathcal{I}^{un}$ in static games, any μ_i is nearly stationary.

It is immediate from Definition 3 that a strategy can have at most one limit strategy. However, a strategy need not have any limit strategy. For a counterexample, consider a game with a single player receiving some signal s_i where $s_i-\theta$ has distribution F_i conditional on θ, and suppose the action space given s_i is $\{s_i,s_i+1\}$. Then the following strategy, call it μ_i, has no limit strategy: assign action $a_i=s_i+1$ for all signals $s_i\geq 0$ and action $a_i=s_i$ for all signals $s_i<0$. As $\theta\to+\infty$, $\mu_i^\theta\to\kappa_0$, where we use κ_x to denote the stationary strategy characterized by $a_i(s_i)=s_i+x$ with probability 1 for all $s_i\in\mathbb{R}$. Hence κ_0 is the only candidate for a limit strategy. But as $\theta\to-\infty$, $\mu_i^\theta\to\kappa_1\neq\kappa_0$, and thus μ_i has no limit strategy.

To see that nearly stationary does not imply stationary, modify the example above so that μ_i is characterized by $a_i=s_i+1$ for all $s_i\in[-K,K]$ for some $K>0$, and $a_i=0$ otherwise. We have $\lim_{\theta\to-\infty}\mu_i^\theta=\lim_{\theta\to+\infty}\mu_i^\theta=\kappa_0$, so μ_i is nearly stationary, but it is not stationary.

The next lemma provides a useful property equivalent to the one in Definition 3. It says that a strategy μ_i^* is a limit strategy for μ_i if and only if μ_i is close to μ_i^* for "most" θ, that is for all θ except for some set of finite measure. Recall that λ denotes the Lebesgue measure.

Lemma 2 *For $i \in \mathcal{I}^{inf}$, a strategy μ_i^* is a limit strategy for μ_i if and only if $\lambda\left(\{\theta : d_{P,i}(\mu_i^\theta, \mu_i^*) \geq \eta\}\right) < \infty$ for all $\eta > 0$, where $d_{P,i}$ is the Prokhorov metric.*

The following Lemma shows that the first part of Definition 3 implies the second.

Lemma 3 *If μ_i^* is a limit strategy of μ_i in the sense of Definition 3, then μ_i^* is stationary and thus μ_i is nearly stationary.*

Proof If μ_i^* is not stationary, then there exists θ such that $\mu_i^{*,\theta} \neq \mu_i^*$. Now if μ_i^* is a limit strategy of μ_i, then $\mu_i^{*,\theta}$ is a limit strategy of μ_i^θ. But also note that if μ_i^* is a limit strategy of μ_i, then μ_i^* is a limit strategy of $\mu_i^{\theta'}$ for all $\theta' \in \mathbb{R}$, including $\theta' = \theta$, and thus μ_i^* is a limit strategy of μ_i^θ. Since a strategy can have at most one limit strategy, we have $\mu_i^{*,\theta} = \mu_i^*$, a contradiction. We conclude that μ_i^* is stationary.

For ease of exposition, from now on we require that the game involves no redundant strategies, i.e., strategies which are indistinguishable from other strategies in their payoff implications.

Assumption 3 (Irreducibility) *The game is irreducible in that there are no distinct strategies $\mu_i \neq \mu_i'$ for any player i such that for all $\theta \in \mathbb{R}$ and all profiles μ_{-i} of stationary strategies for the rivals, $u_i(\mu_i, \mu_{-i}, \theta) = u_i(\mu_i', \mu_{-i}, \theta)$.*

3.3 The Diffuse Prior

The *informal* concept of diffuse prior has two key properties. The first property is that all real numbers are in the support of the prior. The second property is uniformity — all points are weighted equally. Hence, we define a sequence of proper measures to be *diffusing* if these properties hold in the limit — that is, sufficiently far into the sequence, the properties hold arbitrarily closely. The following definition formalizes this idea.

Definition 4 (Diffusing Sequence) Consider a sequence $(\mathbb{P}_n)_{n\in\mathbb{N}}$ of Borel probability measures on $\mathbb{R}$. We say that this sequence is *diffusing* if for any set W with $\lambda(W) \in (0, \infty)$ and any $\eta > 0$, there exists $N \in \mathbb{N}$ such that for all $n \geq N$,

- $\mathbb{P}_n(W) > 0$, and
- for all measurable $Y \subseteq W$, $\left|\frac{\mathbb{P}_n(Y)}{\mathbb{P}_n(W)} - \frac{\lambda(Y)}{\lambda(W)}\right| < \eta$.

From the definition, we can establish the following property of diffusing sequences as a result.

Lemma 4 *If $(\mathbb{P}_n)$ is diffusing, then for any set $E \subset \mathbb{R}$ with $\lambda(E) < \infty$, $\lim_{n\to\infty} \mathbb{P}_n(E) = 0$.*

Proof Consider any set $E \subset \mathbb{R}$ with $\lambda(E) < \infty$ and any diffusing sequence $(\mathbb{P}_n)$. We establish the result for $\lambda(E) > 0$, from which the result for $\lambda(E) = 0$ immediately follows. Choose any arbitrarily small $\eta > 0$. Choose $M > 1$ sufficiently large that (i) $\frac{1}{M} < \eta$ and (ii) $\lambda(E \setminus [-M, M]) < \eta\lambda(E)$. Let $W_M := [-M, M]$ and $W_{M^2} := [-M^2, M^2]$. By applying Definition 4 twice, first with the set E playing the role of W in the definition and again with W_{M^2} playing the role of W, there exists K such that for all $n \geq K$,

$$\mathbb{P}_n(W_{M^2}) > 0$$

$$\mathbb{P}_n(E) > 0$$

$$\left|\frac{\mathbb{P}_n(E \setminus W_M)}{\mathbb{P}_n(E)} - \frac{\lambda(E \setminus W_M)}{\lambda(E)}\right| < \eta \tag{2}$$

$$\left|\frac{\mathbb{P}_n(W_M)}{\mathbb{P}_n(W_{M^2})} - \frac{\lambda(W_M)}{\lambda(W_{M^2})}\right| < \eta. \tag{3}$$

In (2), recall that by construction $\frac{\lambda(E \setminus W_M)}{\lambda(E)} < \eta$ and thus rearranging (2) yields

$$\mathbb{P}_n(E \setminus W_M) < 2\eta\mathbb{P}_n(E) < 2\eta. \tag{4}$$

In (3), we have $\frac{\lambda(W_M)}{\lambda(W_{M^2})} = \frac{1}{M} < \eta$ and thus

$$\mathbb{P}_n(W_M) < 2\eta\mathbb{P}_n(W_{M^2}) < 2\eta. \tag{5}$$

Adding (4) and (5) then yields

$$\mathbb{P}_n(E) \leq \mathbb{P}_n(E \setminus W_M) + \mathbb{P}_n(W_M) < 4\eta.$$

Since η is arbitrary, we have $\mathbb{P}_n(E) \to 0$.

For illustrative purposes, we highlight two specific diffusing sequences (see Figure 1). As one would expect, flattening sequences of the uniform distribution or normal distribution are diffusing according to our definition.

Example 2 Both $(\mathbb{P}_n^1)$ given by the density $\frac{1}{2n}\mathbb{1}_{[-n,n]}$ and $(\mathbb{P}_n^2)$ given by $N(0,n)$ are diffusing.

Proof If $\lambda(W) \in (0,\infty)$, then $W \cap [-n,n] \neq \emptyset$ and thus $\mathbb{P}_n^1(W) > 0$ for sufficiently large n. Moreover, for any measurable $Y \subseteq W$, $\left|\frac{\mathbb{P}_n^1(Y)}{\mathbb{P}_n^1(W)} - \frac{\lambda(Y)}{\lambda(W)}\right| = \left|\frac{\int_{y\in Y}\frac{1}{2n}\mathbb{1}_{[-n,n]}(y)dy}{\int_{w\in W}\frac{1}{2n}\mathbb{1}_{[-n,n]}(w)dw} - \frac{\lambda(Y)}{\lambda(W)}\right| = \left|\frac{\int_{y\in Y}\mathbb{1}_{[-n,n]}(y)dy}{\int_{w\in W}\mathbb{1}_{[-n,n]}(w)dw} - \frac{\lambda(Y)}{\lambda(W)}\right| \to 0$ by dominated convergence. Now $\mathbb{P}_n^2(W) > 0$ for all $n \in \mathbb{N}$. Using a similar argument, for any $Y \subseteq W$, $\left|\frac{\mathbb{P}_n^2(Y)}{\mathbb{P}_n^2(W)} - \frac{\lambda(Y)}{\lambda(W)}\right| = \left|\frac{\int_{y\in Y}\frac{1}{\sqrt{2\pi n}}e^{-y^2/(2n)}dy}{\int_{w\in W}\frac{1}{\sqrt{2\pi n}}e^{-w^2/(2n)}dw} - \frac{\lambda(Y)}{\lambda(W)}\right| = \left|\frac{\int_{y\in Y}e^{-y^2/(2n)}dy}{\int_{w\in W}e^{-w^2/(2n)}dw} - \frac{\lambda(Y)}{\lambda(W)}\right| \to 0.$

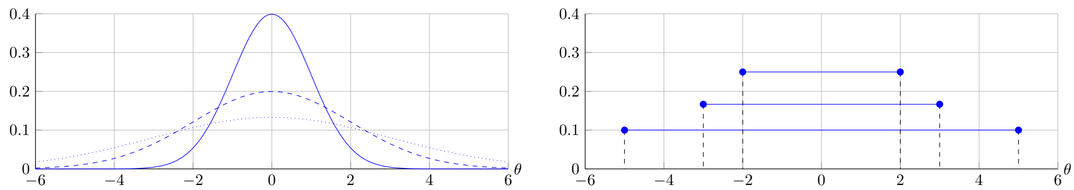

(a) Flattening Normal Distributions

(b) Flattening Uniform Distributions

Fig. 1: Examples of Diffusing Sequences

3.4 Admissibility

Recall the example from Section 2, where we asked which functions of θ are "admissible" in the sense that they yield consistent limits when integrated along certain kinds of sequences of distributions. We suggested that admissible functions are nearly constant in a particular way: there exists some constant such that for any $\epsilon > 0$, the set of θ on which the function deviates from that constant by more than ϵ has finite measure. As we show in our main result, the spirit of this example extends to admissibility of strategies. Since players only observe noisy signals of θ, their strategies are not directly functions of θ but functions of their private signals. Nonetheless, having well-defined expected payoffs will place restrictions on players' strategies, similar to how admissible functions are nearly constant.

Definition 5 (Admissibility) A class $\mathcal{M}^0 \subset \mathcal{M}$ of strategy profiles is said to be *admissible* if for any profile $(\mu_1, \mu_2, ..., \mu_N)$ of strategies in $\mathcal{M}^0$, there exists a vector $u^* \in \mathbb{R}^N$ such that for any diffusing sequence $(\mathbb{P}_n)$ and all $i \in \mathcal{I}$, $\lim_{n\to\infty}\int_{\theta\in\mathbb{R}} u_i(\mu_i, \mu_{-i}, \theta)\, d\mathbb{P}_n(\theta) = u_i^*$.

Let $\mathcal{K}_i$ denote the class of stationary strategies for player i. Note that for profiles $\mu = (\mu_i, \mu_{-i})$ of stationary strategies, the mapping $\theta \mapsto u_i(\mu, \theta)$ is constant, and hence the sequence of integrals in Definition 5 trivially converges to the same limit for every diffusing sequence. Thus the class of profiles of stationary strategies is admissible.

Given the definition of admissibility, we can define (ex ante) payoff equivalence between two strategies.

Definition 6 (Payoff Equivalence) Let $\mathcal{M}^0$ be a class of admissible strategy profiles. Two strategies μ_i and μ_i' are said to be payoff equivalent for player i if for all strategy profiles μ_{-i} of the remaining players such that $(\mu_i, \mu_{-i}), (\mu_i', \mu_{-i}) \in \mathcal{M}^0$, player i gets the same payoff under μ_i as under μ_i': for all diffusing sequences $(\mathbb{P}_n)$,

$$\lim_{n\to\infty} \int_{\theta\in\mathbb{R}} u_i(\mu_i, \mu_{-i}, \theta)\, d\mathbb{P}_n(\theta) = \lim_{n\to\infty} \int_{\theta\in\mathbb{R}} u_i(\mu_i', \mu_{-i}, \theta)\, d\mathbb{P}_n(\theta).$$

We now present our main result for single-stage games. For the necessity direction, our proof technique makes use of the assumption that the class of strategy profiles contains stationary strategies. The reason is that when we consider payoffs conditional on various θ, stationarity of the rivals' strategies (together with stationarity of signals and payoffs) means that an informed player i's payoff is changing *as if* one is holding θ fixed and changing the actual strategy of player i. By irreducibility, this means that if i's strategy is not nearly stationarity, then i's payoff conditional on θ varies significantly with θ, which we then show is inconsistent with admissibility.

Theorem 1 *Suppose the game Γ satisfies Assumptions 1-3. Then*

- *(Sufficiency) The class of profiles of nearly stationary strategies is admissible.*
- *(Necessity) If $\mathcal{M}^0 = \times_{i\in\mathcal{I}}\mathcal{M}_i^0$ is admissible and $\mathcal{K}_i \subseteq \mathcal{M}_i^0$ for all i, then every $\mu_i \in \mathcal{M}_i^0$ is nearly stationary and payoff equivalent to some stationary strategy.*

4 Multistage Games

We now extend the analysis to a class of multistage games. The game has stages indexed by $t = 1, 2, \ldots, \tau$ for $\tau \in \mathbb{N}$. In each stage t, some nonempty subset of players choose actions $a_i^t \in \mathbb{R}$; we use $\mathcal{I}(t)$ to denote this subset. As in the static model, a subset of players $\mathcal{I}^{inf}$ are informed and obtain a private signal s_i with the same properties as in the static game, while the remaining players receive no such signal. Each player acts in exactly one stage, denoted $T(i)$. We assume that informed players act in stage 1, while all uninformed players act in later stages. We assume that uninformed players have a finite action set. Players observe other players' past actions perfectly. The public history at stage $t > 1$ is a record of all players' past actions through stage $t-1$ and it is denoted h^t. An informed player's history is simply that player's initial signal s_i. To economize on notation, we define the history observed by player $i \in \mathcal{I}(t)$ by $h_i^t := h^t$ when $t > 1$ (and thus $i \in \mathcal{I}^{un}$) and $h_i^t := s_i$ when $t = 1$ (and thus $i \in \mathcal{I}^{inf}$). The set of available actions for $i \in \mathcal{I}$ moving in stage t is $A_i^t(h_i^t)$.

A pure strategy specifies for each player's observed history h_i^t an action in $A_i^t(h_i^t)$. Given a realization of θ and sequence of action profies taken in each stage, players receive payoffs $u_i(a_i, a_{-i}, \theta)$, where a_{-i} is defined as before. In Definition 7, we extend the definition of distributional strategies from Section 3.

In order to extend Theorem 1, we must first adapt some other concepts from Section 3. As in Section 3, payoffs are assumed stationary in that they are invariant to shifts in θ and all informed players' actions by the same constant. We adapt Assumption 1 to the following:

Assumption 4 *For all $i \in \mathcal{I}(1)$, $A_i^1(h_i^1)$ is compact and satisfies $A_i^1(s_i) = A_i^1(0) + s_i$. For all $t > 1$ and all $i \in \mathcal{I}(t)$, $A_i^t(h_i^t)$ is finite and independent of h_i^t.*

Given Assumption 4, we use A_i^t to denote the action sets available to uninformed players.

Given a realization of θ and sequence of action profiles taken in each stage, players receive payoffs $u_i(a_i, a_{-i}, \theta)$ for the game, where a_{-i} is defined as before. As before, we assume that the payoff functions u_i are continuous in all arguments, and we assume they are stationary in the same sense as in Section 3.

As in the static game, informed players' only information is their private signal, so their strategies are defined exactly as in the static game. For uninformed players, their information is the past actions of other players. For $t > 1$, let $H^t := (\times_{i\in\mathcal{I}^{inf}} X_i) \times \left(\times_{i\in\mathcal{I}^{un}:T(i)<t} A_i^{T(i)}\right)$ denote the set of public histories at stage t. We formulate the strategies for players acting in stage $t > 1$ as joint distributions over H^t and their stage t actions, fixing a marginal distribution over H^t. Any version of the conditional distribution specifies behavior after a given history; under regularity conditions discussed below, any two versions of the conditional distribution are outcome-equivalent.

Definition 7 [Strategies — Multistage Version] A strategy for player $i \in \mathcal{I}^{inf}$ is a probability measure μ_i on $S_i \times X_i$ satisfying the properties of Definition 1. A strategy for player $i \in \mathcal{I}^{un}$ acting in stage t is a probability measure μ_i on $H^t \times A_i^t$ with the following property: the random variables induced by the marginal distributions over the dimensions for actions of players acting earlier are mutually independent and have distribution $N(0,1)$ for informed players and (discrete) uniform distribution for uninformed players.

For an uninformed player i, the distributional strategy μ_i includes dimensions for all the other players' past actions. This results in some arbitrary detail which serves a technical purpose only. The strategically meaningful information from μ_i is in the distributions on the A_i^t dimension conditional on each history. The marginal distributions for the past actions are *not* strategically meaningful, so for concreteness we normalize those to the normal distribution for past actions by informed players and discrete uniform distributions over finite sets for past actions by uninformed players; we wish to emphasize that μ_i does not contain any conjecture about other players' behavior. Part (iii) of Assumption 5 below ensures that all actions by informed players occur with probability zero (such as when all informed players play stationary strategies, due to the signal noise), and hence an uninformed player i's behavior is well-defined, as all versions of the uninformed players' strategies are outcome-equivalent.

Hence, there is an equivalence between distributional strategies under our formulation and behavioral strategies, given the regularity condition mentioned above. Specifically, for every distributional strategy, there exists an outcome-equivalent behavioral strategy, and conversely, for every behavioral strategy, there exists an outcome-equivalent distributional strategy. A distributional strategy specifies behavior at each on-path decision node via the conditional distributions over that player's actions, conditioning on the history; these conditional distributions and hence the player's behavior are uniquely determined except on a set of measure zero. Given a behavioral strategy, one can construct a distributional strategy by specifying the marginal distributions on earlier actions as in Definition 7, and by defining the conditional distributions as those under the behavioral strategy.

Each player has a set of available strategies denoted $\mathcal{M}_i$; as in Section 3, it is metrized by the Prokhorov distance. In multistage games, some technical issues not present in static games arise. For instance, small changes in one player's strategy can lead to discontinuous changes in the induced distribution over outcomes. To avoid such issues, we assume that $\mathcal{M} := \times_{i\in\mathcal{I}}\mathcal{M}_i$ satisfies some regularity conditions (Assumption 5). Part (iii) ensures that each zero-measure set of actions occurs with zero probability, as described above. Part (v) of the definition adapts irreducibility from Assumption 3. As in the single-stage case, given any θ realization, a profile of distributional strategies (μ_i, μ_{-i}) induces a distribution over payoffs $u_i(a_i, a_{-i}, \theta)$. We again abuse notation and let $u_i(\mu_i, \mu_{-i}, \theta)$ denote ex interim payoffs conditional on θ.

Assumption 5 *The class $\mathcal{M} = \times_{i\in\mathcal{I}}\mathcal{M}_i$ of profiles of available strategies satisfies the following conditions: (i) $\mathcal{M}_i$ closed for all $i \in \mathcal{I}$, (ii) $\mu_i \in \mathcal{M}_i$ implies $\mu_i^\theta \in \mathcal{M}_i$ for all $\theta \in \mathbb{R}$ and $i \in \mathcal{I}$, (iii) $\mu_i(S_i \times Y) = 0$ for all $i \in \mathcal{I}$ and $Y \subset X_i$ such that $\lambda(Y) = 0$, (iv) for all $i \in \mathcal{I}$, $\mu \mapsto u_i(\mu, 0)$ is continuous on $\mathcal{M}$, and (v) for all $i \in \mathcal{I}$ and strategies $\mu_i \neq \mu_i'$ in $\mathcal{M}_i$, there exists $\theta_0 \in \mathbb{R}$ and a profile $\mu_{-i} \in \mathcal{M}_{-i}$ of stationary strategies for the other players such that $u_i(\mu_i, \mu_{-i}, \theta_0) \neq u_i(\mu_i', \mu_{-i}, \theta_0)$.*

Given Definition 7 and Assumption 4, Lemma 1 can be applied to the space of measures satisfying Definition 7, and hence each $\mathcal{M}_i$, as a closed subspace by Assumption 5, is compact. By continuity of interim payoffs in strategies, interim payoffs are bounded.

Recentered strategies: The recentering function for an informed players is exactly as in the static model. For uninformed players, let $h_i^{t,x}$ denote the history h_i^t modified by adding x to all informed players' (past) actions. We define the recentered strategy μ_i^x as the strategy which satisfies, for all $h_i^t \in H^t$ except on a set of measure zero and all $Y \subseteq A_i^t(h_i^t)$, $\mu_i^x(Y|h_i^t) = \mu_i(Y - x|h_i^{t,-x})$. As in the single-stage case, for $x = 0$, we have $\mu_i^0 = \mu_i$. For strategy profiles, we define μ^x as the profile in which each player plays μ_i^x. Under the generalized recentering function defined above, we can now adapt Definition 2 to multistage games.

Definition 8 A strategy μ_i is stationary if for all $x \in \mathbb{R}$, $\mu_i = \mu_i^x$.

We can import several statements from Section 3 with no change in notation and therefore there is no need to reproduce them here. In particular, we maintain the following definitions: (i) the definition of limit strategies and nearly stationary strategies (Definition 3), (ii) the definition of a diffusing sequence (Definition 4), (iii) the definition of admissible strategies (Definition 5), and (iv) the definition of payoff equivalence (Definition 6). We define $\mathcal{K}_i$ as the class of stationary strategies in $\mathcal{M}_i$.

We now extend Theorem 1 to multistage games.

Theorem 2 *Suppose Assumptions 4 and 5 are satisfied.*

– *(Sufficiency) The class of profiles of nearly stationary strategies is admissible.*
– *(Necessity) If* $\mathcal{M}^0 = \times_{i\in\mathcal{I}}\mathcal{M}_i^0$ *is admissible and* $\mathcal{K}_i \subseteq \mathcal{M}_i^0$ *for all* i*, then every* $\mu_i \in \mathcal{M}_i^0$ *is nearly stationary and payoff equivalent to some stationary strategy.*

5 Applications

In this section, we apply our single-stage model to a beauty contest game and our multistage model to a delegation (sender-receiver) game.

5.1 Beauty Contests

An application of the single-stage version of our model is an adaptation of the beauty contest model of Morris and Shin (2002). We assume that there is a finite number $N \geq 2$ of agents $i \in \{1, 2, \ldots, N\}$ and an underlying state of the world θ drawn from a diffuse prior distribution. Each agent has a bias parameter b_i and receives a private signal $s_i = \theta + \epsilon_i$, where $\epsilon_i \sim N(0, \sigma_\epsilon^2)$ and where the ϵ_i and θ are mutually independent. Each agent chooses an action $a_i \in [s_i - M, s_i + M]$ for some large constant $M > 0$ and receives a payoff

$$u(a_i, a_{-i}, \theta) = -\lambda(a_i - \theta - b_i)^2 - (1 - \lambda)(a_i - \bar{a})^2,$$

where $\lambda \in (0, 1)$ is a constant and $\bar{a} =:= \frac{\sum_j a_j}{N}$ is the average of all players' actions, including player i. In other words, players want to choose actions that balance their idiosyncratic preferences with a desire to match other agents' actions.

The game described here has stationary payoffs and satisfies Assumptions 1-3, and thus Theorem 1 applies. Hence, in order to guarantee well-defined ex ante expected payoffs, the analyst must restrict strategy sets to nearly stationary strategies. In the game we have described, stationary strategies are mixtures over constant markups: each player chooses a (possibly random) markup $k_i \in [-M, M]$ and, given any s_i, plays the action $a_i(s_i) = s_i + k_i$.

If we restrict players to stationary strategies, there is a unique Bayesian Nash Equilibrium of the game with diffuse prior, in which each player plays a pure strategy characterized by $a_i(s_i) = s_i + k_i$, where

$$k_i = \lambda b_i + (1 - \lambda)\frac{\sum_j b_j}{N}. \tag{6}$$

Each player's signal s_i is that player's posterior mean belief about the state θ and about all other agent's signals s_j for $j \neq i$, and thus the mean action of each other player is $s_i + k_j$. Each player's payoff is a concave function and thus has a unique maximizer given the strategies of the other players. It follows that any equilibrium is in pure strategies. Equation (6) says that each player's equilibrium markup is a convex combination of all players' biases, and the weight assigned to one's own bias is an increasing function of λ. In the extreme case $\lambda = 1$, each player cares only about his idiosyncratic preferences, and optimally sets $k_i = b_i$. For the other extreme, $\lambda = 0$, the game is a pure coordination game and multiple equilibria exist, but in the limit as $\lambda \to 0$, $k_i \to \frac{\sum_j b_j}{N}$; with agents coordinating on the average bias of all agents.

5.2 Delegation

As a direct application of our model, we can also consider sender-receiver games. To illustrate, consider games with two stages: one stage in which n senders simultaneously choose actions after observing private signals, and a second stage in which the receiver observes the senders' actions and chooses an action. Depending on the application in mind, the receiver's action could be interpreted as a continuous variable (as in games of cheap talk) or as a sender's identity (as in games of delegation with multiple experts).

In the case of delegation, the receiver's available action choices may depend on the actions chosen by the senders. For example, in a companion paper by Ambrus et al. (2019), the senders are two experts who propose action choices, and the receiver must choose one of them. All players have quadratic loss functions; senders have biases b_i relative to the receiver. The game unfolds as follows:

1. A state of the world $\theta \in \mathbb{R}$ is realized, drawn from a diffuse prior distribution.

2. Senders (experts) $i = 1, 2$ receive conditionally i.i.d. signals $s_i \sim N(\theta, \sigma^2)$.
3. Senders choose markups $k_i \in [-M, M]$ where M is some large constant.
4. The receiver (principal or decision maker, labeled player 3) observes $a_i = s_i + k_i$ for $i \in \{1, 2\}$ and chooses the expert whose action will be implemented. Specifically, the receiver chooses an action j from $\{1, 2\}$.
5. Payoffs are realized: receiver gets $-(s_j + k_j - \theta)^2$, sender j gets $-(s_j + k_j - \theta - b_j)^2$, sender $i \neq j$ gets $-(s_j + k_j - \theta - b_i)^2$.

A stationary strategy for the receiver must have the property that $C(a_1, a_2) = C(a_1 + x, a_2 + x)$ for all x. This includes mixtures of the following fundamental strategies: (i) always choose the higher offer, (ii) always choose the lower offer, (iii) always choose the offer from sender 1, and (iv) always choose the offer from sender 2. Ambrus et al. (2019) show that these are essentially the only possible best responses of the principal to a pair of stationary strategies of the experts.

Suppose the receiver is restricted to the four pure strategies above, and denote this strategy space by $\mathcal{R}$. We show that this strategy space alone is enough to establish irreducibility for the senders. Specifically, we show that any class $\mathcal{M} = \mathcal{M}_1 \times \mathcal{M}_2 \times \mathcal{R}$ of strategy profiles that satisfies parts (i)-(iv) of Assumption 5 must also satisfy part (v), and hence Theorem 2 applies: all admissible strategies are nearly stationary, and in particular, are nearly "constant markup strategies" — that is, each sender's strategy is nearly a stationary strategy in which there is a single "markup" distribution H_i over $[-M, M]$ and actions are simply s_i plus the draw from H_i.

Proposition 1 *In the stationary game of Ambrus et al. (2019), suppose that strategy spaces $\mathcal{M}_1$ and $\mathcal{M}_2$ are such that the class of strategy profiles $\mathcal{M} = \mathcal{M}_1 \times \mathcal{M}_2 \times \mathcal{R}$ satisfies parts (i)-(iv) of Assumption 5. Then part (v) is also satisfied, and Theorem 2 applies to the game.*

6 Alternative Approaches to Nonstationary Strategies

In this section, we consider alternative approaches to handling ex ante payoffs under the diffuse prior.[11]

First, in some cases it is possible to make comparisons across strategy profiles — for example, when an individual player considers a deviation — without requiring that ex ante payoffs be well-defined before and after the deviation. For instance, consider a single-stage game and suppose that all players are playing stationary strategies which are mutual best responses within the class of stationary strategies. To verify that such a strategy profile is a Bayesian Nash equilibrium, one would like to verify that no player can benefit by deviating, including deviations to nonstationary strategies. While nonstationary strategies (except nearly stationary strategies) do not yield well-defined ex ante payoffs, the limit suprema of payoffs are well-defined; optimality implies these must be bounded above by the (well-defined) payoff of the original, stationary strategy. In this sense, deviations can be ruled out. Hence, while our approach would already rule out such deviations on the basis of admissibility, an equilibrium concept could be defined using a stronger notion of best response which allows deviations outside the admissible class.

Second, in some cases it is possible to obtain well-defined expected payoffs from profiles of strategies which are not nearly stationary. Our admissibility criterion has imposed the inclusion of stationary strategies in an admissible class; we consider this inclusion to be a natural requirement for stationary games. However, dropping this requirement allows one to construct smaller admissible classes. Consider the following example: two players $i = 1, 2$ obtain perfect signals $s_i = \theta$ and choose actions $a_i \in [s_i, s_i+1]$. Their payoffs are $u_i(a_1, a_2, \theta) = -(a_1 - a_2)^2(a_i - \theta)^2$. Consider the (singleton) class of strategy profiles $\{(a_1^*, a_2^*)\}$ where $a_i^*(s_i) := s_i + \mathbb{1}_{s_i \geq 0}$. This class is clearly admissible, but the a_i^* are neither stationary nor nearly stationary. On the other hand, the strategy a_2^* does not yield a well-defined payoff against the stationary strategy $\tilde{a}_1$ defined by $\tilde{a}_1(s_1) = s_1$; player 2's ex post payoff would be 0 when $\theta < 0$ but it would be -1 when $\theta \geq 0$.

Third, there may be a "natural" labeling of the state space, or a natural choice of the origin in some games, such as a status quo. In such games, strategies such as $a_i^*(s_i)$ above may be quite reasonable, and one might wish to find an admissible class which includes these. To that end, one could relax our requirement that *all* diffusing sequences of priors yield the same expected payoffs in the limit and instead focus on a restricted class of sequences of priors, or a particular sequence of priors such as $U([-k, k])$ for $k = 1, 2, \ldots$. Note that even a particular sequence of priors such as this one still disciplines the strategies. For example, consider perfect signals $s_i = \theta$ and payoffs $u_i(a_i, \theta) = -(a_i - \theta)^2$, and define strategies by $a_i(s_i) = s_i + 1$ on intervals of the form $[-2^{2n+1}, -2^{2n}]$ or $[2^{2n}, 2^{2n+1}]$ and $a_i(s_i) = s_i$ elsewhere. The

[11] The authors thank an anonymous referee for several observations upon which this section is based.

expected payoff converges to $-2/3$ along a subsequence $U([-2^{2n+1}, 2^{2n+1}])$ but converges to $-1/3$ along a subsequence $U([-2^{2n}, 2^{2n}])$. On the other hand, the strategy $a_i^*(s_i)$ from above, which is not nearly stationary, yields a well-defined expected payoff of $1/2$ in this case. Hence, restriction to a smaller class of sequences of priors appropriate for a given problem can admit new strategy profiles while continuing to exclude others.

7 Conclusion

We have presented a formal method for defining ex ante payoffs in games with diffuse prior. The key features of the diffuse prior can be captured using a limit construction, in which sequences of proper priors exhibit these properties in a limiting sense. Under our construction, stationary strategies admit well-defined payoffs in stationary games, and conversely, all strategies admitting well-defined payoffs are nearly stationary in a precise sense.

Our methodology can be readily extended in several directions. Although we have considered a one-dimensional state of the world θ, in many applications, there is uncertainty over multiple dimensions. One could model a diffuse prior over a multi-dimensional state by generalizing our notion of a diffusing sequence. Another, related direction would be to consider renewed uncertainty; in a multistage game, each stage t might introduce more uncertainty through the realization of a state θ_t. We leave detailed exploration of these directions to future work.

Although we have considered exogenous uncertainty over the state θ, our construction could also be used to allow a player to choose the diffuse prior as a mixed strategy. That is, suppose an additional player $i = 0$ is introduced who chooses an action θ at the beginning of the game. Under our construction, this player can play the diffuse prior as a mixed strategy with a well-defined payoff.

A Proofs

Proof (Proof of Lemma 1.) We prove the result for informed players; the arguments for uninformed players are essentially a subset of those here. Note that $\mathcal{M}_i$ is a metric space, as it is a subspace of the metric space (with the Prokhorov metric) consisting of measures over the complete and separable metric space $(S_i \times X_i, d_{S_i \times X_i})$ where $d_{S_i \times X_i}((s_i, x_i), (s_i', x_i')) = \frac{1}{2}\max\{|s_i - s_i'|, |x_i - x_i'|\}$.[12] To show that $\mathcal{M}_i$ is compact it suffices to show that it is relatively compact and that it is closed. Below we show that $\mathcal{M}_i$ is relatively compact; to verify that $\mathcal{M}_i$ is also closed is straightforward and only requires the tedious steps of showing that convergent sequences of strategies converge to limits which satisfy the key properties of Definition 1.

Given any $\eta > 0$, pick any compact $T \subset S_i$ such that $\phi(T) > 1-\eta$. Define $Z_i \subset S_i \times X_i$ by $Z_i := \{(s_i, x_i) : x_i \in A_i(s_i)\}$. Now Z_i is compact as a consequence of Assumption 1, and it satisfies $\mu_i(Z_i) = \phi(T) > 1 - \eta$. Hence, $\mathcal{M}_i$ is tight and by Prokhorov's Theorem,[13] it is relatively compact. Together with closedness, we conclude that $\mathcal{M}_i$ is compact.

Before proving Lemma 2, we first state and prove Lemmas 5 and 6; this does not involve any circularity.

Lemma 5 *For all $i \in \mathcal{I}$ and all strategies $\mu_i \in \mathcal{M}_i$, the map $\theta \mapsto \mu_i^\theta$ from $\mathbb{R}$ to $\mathcal{M}_i$ is uniformly continuous.*

Proof The result is trivial for uninformed players, since μ_i^θ is constant in θ. Hence, we prove the result for informed players. We first specify a few preliminaries. We recall the space $S_i \times X_i$ is metrized by $d_{S_i \times X_i}$ defined in the proof of Lemma 1. For any subset $Y \subseteq S_i \times X_i$ and $\eta > 0$, let $Y^\eta := \bigcup_{z \in Y} N_\eta(z)$, the union of all η-neighborhoods (under the metric $d_{S_i \times X_i}$) centered at points in Y. It follows from these definitions that for any $\eta \in \mathbb{R}$ (possibly negative), $Y + \eta \subseteq Y^{|\eta|}$, where by our notational convention $Y + \eta$ denotes the translation of the set Y by η (with respect to the standard metric on $\mathbb{R}$) in all dimensions. The space $\mathcal{M}_i$ is metrized by the Prokhorov metric, $d_{P,i}(\mu, \hat{\mu}) := \inf\{\eta > 0 : \mu(Y) \leq \hat{\mu}(Y^\eta) + \eta \text{ and } \hat{\mu}(Y) \leq \mu(Y^\eta) + \eta \text{ for all } Y \subseteq S_i \times X_i\}$.

For the proof, we show that for all $\eta > 0$, if $|\theta - \theta'| < \eta$, then $d_{P,i}(\mu^\theta, \mu^{\theta'}) < \eta$. For all $Y \subseteq S_i \times X_i$ and $s_i \in S_i$, define $Y(s_i) := \{x_i \in X_i : (s_i, x_i) \in Y\}$. We have

$$\begin{aligned}\mu^\theta(Y) = \int_{s_i \in S_i} \mu^\theta(Y(s_i)|s_i)dG(s_i) &= \int_{s_i \in S_i} \mu^{\theta'}(Y(s_i) + \theta' - \theta|s_i)dG(s_i) \\ &\leq \int_{s_i \in S_i} \mu^{\theta'}(Y^{|\theta'-\theta|}(s_i)|s_i)dG(s_i) = \mu^{\theta'}(Y^{|\theta'-\theta|}) < \mu^{\theta'}(Y^\eta) + \eta,\end{aligned}$$

and by a symmetric argument $\mu^{\theta'}(Y) < \mu^\theta(Y^\eta) + \eta$. By the definition of $d_{P,i}$, $d_{P,i}(\mu^\theta, \mu^{\theta'}) < \eta$.

Lemma 6 *For all $i \in \mathcal{I}$, $u_i(\mu, \theta)$ is uniformly continuous in μ_j for each $j \in \mathcal{I}$, and in θ.*

[12] The factor of 1/2 here is not necessary, but it is useful in the proof of Lemma 5, so we adopt it here for consistency.
[13] See, Billingsley (2009, Theorem 5.1).

Proof To establish that $u_i(\mu,\theta)$ is continuous in each μ_j, for all strategy profiles μ and all θ, define $\tilde{\mu}$ is the product measure over $Z := (\times_{i\in\mathcal{I}} X_i)\times S$ induced by the strategy profile μ, modified so that the marginal distribution over S_i has density $f_i(s_i-\theta)$. We can then write $u(\mu,\theta)=\int_{(a,s)\in Z} u_i(a,\theta)d\tilde{\mu}(a,s)$. Let $\epsilon>0$ and $\mu\in\mathcal{M}$ be arbitrary. Fix any compact set $\overline{S}\subseteq\mathbb{R}^{|\mathcal{I}^{inf}|}$ of the form $\overline{S}=\times_{i\in\mathcal{I}^{inf}}[\underline{s}_i,\overline{s}_i]$ large enough so that $\int_{\hat{s}\notin\overline{S}}\sup_{a\in A(\hat{s}+\theta)}|u(a_i,a_{-i},\theta)|dF(\hat{s})<\epsilon$, where we have used Assumption 2. Let $\overline{Z}:=(\times_{i\in\mathcal{I}}X_i)\times\overline{S}$, which is compact and has zero-measure boundary under $\tilde{\mu}$. Suppose a sequence $(\mu^n)_{n\in\mathbb{N}}$ converges to μ; then $\tilde{\mu}^n\to\tilde{\mu}$, and thus for sufficiently large N, for all $n\geq N$,

$$\begin{aligned}|\int_{(a,s)\in Z}u_i(a,\theta)d\tilde{\mu}^n(a,s)-\int_{(a,s)\in Z}u_i(a,\theta)d\tilde{\mu}(a,s)| &\leq |\int_{(a,s)\in\overline{Z}}u_i(a,\theta)d\tilde{\mu}^n(a,s)-\int_{(a,s)\in\overline{Z}}u_i(a,\theta)d\tilde{\mu}(a,s)|\\ &+|\int_{(a,s)\in Z\setminus\overline{Z}}u_i(a,\theta)d\tilde{\mu}^n(a,s)|+|\int_{(a,s)\in Z\setminus\overline{Z}}u_i(a,\theta)d\tilde{\mu}(a,s)|\\ &<3\epsilon.\end{aligned}$$

Hence $u_i(\mu,\theta)$ is continuous in each μ_j, and since the μ_j lie in a compact domain, this continuity is uniform. Next, since $u_i(\mu,\theta)=u_i(\mu^{-\theta},0)$ and by Lemma 5 each $\mu_j^{-\theta}$ term is uniformly continuous in θ, we have uniform continuity in θ.

Proof (Proof of Lemma 2.) For the "only if" direction, note that by the definition of a limit, $\lim_{\theta\to+\infty}\mu_i^\theta=\lim_{\theta\to-\infty}\mu_i^\theta=\mu_i^*$ implies that for all $\eta>0$, there exists $M>0$ such that for all $\theta\leq -M$ and all $\theta\geq M$, $d_{P,i}(\mu_i^\theta,\mu_i^*)<\eta$. Hence $\{\theta: d_{P,i}(\mu_i^\theta,\mu_i^*)\geq\eta\}\subseteq[-M,M]$ and thus $\lambda\left(\{\theta: d_{P,i}(\mu_i^\theta,\mu_i^*)\geq\eta\}\right)<\infty$. For the "if" direction, we prove the contrapositive. Suppose that some arbitrary strategy μ_i^* is not a limit strategy for μ_i. Then given any $\eta>0$, it is possible to construct a sequence $(\theta_j)_{j\in\mathbb{N}}$ such that (i) for all $j,k\in\mathbb{N}$ with $j\neq k$, $|\theta_j-\theta_k|>1$ and (ii) for all $j\in\mathbb{N}$, $d_{P,i}(\mu_i^{\theta_j},\mu_i^*)\geq 2\eta$. Since the map $\theta\mapsto\mu_i^\theta$ is uniformly continuous (Lemma 5), there exists $\delta>0$ such that whenever $|\theta-\theta'|<\delta$, $d_{P,i}(\mu_i^\theta,\mu_i^{\theta'})<\eta$. It follows that for all $k\in\mathbb{N}$ and all $\theta'\in(\theta_k-\delta,\theta_k+\delta)$, $d_{P,i}(\mu_i^{\theta'},\mu_i^*)\geq d_{P,i}(\mu^{\theta_k},\mu_i^*)-d_{P,i}(\mu_i^{\theta_k},\mu_i^{\theta'})>2\eta-\eta=\eta$. Hence $\bigcup_{k\in\mathbb{N}}(\theta_k-\delta,\theta_k+\delta)\subseteq\{\theta: d_{P,i}(\mu_i^\theta,\mu_i^*)\geq\eta\}$, and since $\lambda\left(\bigcup_{k\in\mathbb{N}}(\theta_k-\delta,\theta_k+\delta)\right)=\infty$, we have $\lambda\left(\{\theta: d_{P,i}(\mu_i^\theta,\mu_i^*)\geq\eta\}\right)=\infty$, concluding the proof of the contrapositive.

For the proof of the main result, we make use of a weaker property than that of Definition 3.

Lemma 7 *For any $i\in\mathcal{I}$ and strategy $\mu_i\in\mathcal{M}_i$, there exists $\mu_i^*\in\mathcal{M}_i$ with the property that for all $\eta>0$, $\lambda\left(\{\theta\in\mathbb{R}: d_{P,i}(\mu_i^\theta,\mu_i^*)<\eta\}\right)=\infty$, where λ denotes the Lebesgue measure and $d_{P,i}$ denotes the Prokhorov distance defined on $\mathcal{M}_i$. We say that any such μ_i^* is an* attraction *for μ_i.*

Proof We prove a more general result. Let (C,d) be a compact metric space, $B\subseteq\mathbb{R}$ with $\lambda(B)=\infty$, and $\pi: B\to C$ a Lebesgue measurable function. We show there exists $c\in C$ with the property that for all $\eta>0$, $\lambda\left(\{b\in B: d(\pi(b),c)<\eta\}\right)=\infty$. Suppose on the contrary that for each $c\in C$, there exists $\eta_c>0$ such that $\lambda\left(\{b\in B: d(\pi(b),c)<\eta_c\}\right)<\infty$. The collection $\{N_{\eta_c}(c): c\in C\}$ is an open covering of C, and by compactness, it has a finite subcovering denoted $\{N_{\eta_i}(c_i)\}_{i=1}^n$ for some $n\in\mathbb{N}$. It follows that $B\subseteq\cup_{i=1}^n\pi^{-1}(N_{\eta_i}(c_i))$ and thus $\lambda(B)\leq\sum_{i=1}^n\lambda\left(\pi^{-1}(N_{\eta_i}(c_i))\right)<\infty$, a contradiction. Thus there exists c such that the property holds. To conclude, note that this result specializes to the lemma statement by setting $C=\mathcal{M}_i$, $d=d_{P,i}$, $B=\Theta=\mathbb{R}$ and π to be the map $\theta\mapsto\mu_i^\theta$; given the existence of c as above, we set $\mu_i^*=c$.

Lemma 7 says that every strategy has at least one attraction. Note that the example immediately following Definition 3 has two attractions, κ_0 and κ_1, but as shown in that example, it is not nearly stationary.

Proof (Proof of Theorem 1.) We begin with sufficiency.

Sufficiency: Let $\mathcal{M}_i^*$ be the class of nearly stationary strategies for each player i, and consider any profile $(\mu_1,\mu_2,\ldots,\mu_N)$ of such strategies.

By definition, for each $\mu_i\in\mathcal{M}_i^*$, there exists a strategy μ_i^* such that μ_i^* is a limit strategy for μ_i; by Lemma 3, this μ_i^* is stationary and by Lemma 2, for all $\eta>0$,

$$\lambda\left(\{\theta: d_{P,i}(\mu_i^\theta,\mu_i^*)\geq\eta\}\right)<\infty. \tag{7}$$

Next, for all $\eta>0$, define $\Theta_{\geq\eta}=\{\theta\in\mathbb{R}:\max_{i\in\mathcal{I}}d_{P,i}(\mu_i^\theta,\mu_i^*)\geq\eta\}$, which has finite measure by (7) and the fact that there are finitely many players. Let $\mathbb{P}_n$ be any diffusing sequence. For all n, by definition of the recentering function,

$$\int_{\theta\in\mathbb{R}}u_i(\mu_i,\mu_{-i},\theta)\,d\mathbb{P}_n(\theta)=\int_{\theta\in\mathbb{R}}u_i(\mu_i^{-\theta},\mu_{-i}^{-\theta},0)\,d\mathbb{P}_n(\theta). \tag{8}$$

We can write the RHS of (8) as

$$\int_{\theta\in\Theta_{\geq\eta}}u_i(\mu_i^{-\theta},\mu_{-i}^{-\theta},0)\,d\mathbb{P}_n(\theta)+\int_{\theta\in\mathbb{R}\setminus\Theta_{\geq\eta}}u_i(\mu_i^{-\theta},\mu_{-i}^{-\theta},0)\,d\mathbb{P}_n(\theta). \tag{9}$$

In the first term of (9), by Assumption 2, the integrand $u_i(\mu_i^{-\theta},\mu_{-i}^{-\theta},0)$ is bounded in absolute value by some quantity, call it M. For the second term of (9), note that by uniform continuity of ex interim payoffs, for any $\epsilon>0$, there exists $\eta>0$ such that for any profile of strategies $(\tilde{\mu}_1,\tilde{\mu}_2,\ldots,\tilde{\mu}_N)$, if $d_{P,i}(\tilde{\mu}_i,\mu_i^*)<\eta$ for all $i\in\mathcal{I}$, then

$$|u_i(\tilde{\mu}_i,\tilde{\mu}_{-i},0)-u_i(\mu_i^*,\mu_{-i}^*,0)|<\epsilon.$$

Define $u_i^*:=u_i(\mu_i^*,\mu_{-i}^*,0)$ and note that by the definition of $\Theta_{\geq\eta}$, for $\theta\in\mathbb{R}\setminus\Theta_{\geq\eta}$ we have $\max_{i\in\mathcal{I}}d_{P,i}(\mu_i^\theta,\mu_i^*)<\eta$ and thus $|u_i(\mu_i^{-\theta},\mu_{-i}^{-\theta},0)-u_i(\mu_i^*,\mu_{-i}^*,0)|<\epsilon$. Putting these together,

$$\left|\int_{\theta\in\mathbb{R}}u_i(\mu_i,\mu_{-i},\theta)\,d\mathbb{P}_n(\theta)-u_i^*\right|\leq\int_{\theta\in\Theta_{\geq\eta}}|u_i(\mu_i^{-\theta},\mu_{-i}^{-\theta},0)-u_i^*|\,d\mathbb{P}_n(\theta)$$

$$
\begin{aligned}
&+ \int_{\theta \in \mathbb{R} \setminus \Theta_{\geq \eta}} |u_i(\mu_i^{-\theta}, \mu_{-i}^{-\theta}, 0) - u_i^*| \, d\mathbb{P}_n(\theta) \\
&\leq \mathbb{P}_n(\Theta_{\geq \eta})(M + |u_i^*|) + \epsilon \cdot \mathbb{P}_n(\mathbb{R} \setminus \Theta_{\geq \eta}) \\
&\leq \mathbb{P}_n(\Theta_{\geq \eta})(M + |u_i^*|) + \epsilon. \qquad (10)
\end{aligned}
$$

Since $\Theta_{\geq \eta}$ has finite Lebesgue measure, by Lemma 4 there exists K such that for $n \geq K$, $\mathbb{P}_n(\Theta_{\geq \eta}) < \frac{\epsilon}{M+|u_i^*|}$, and thus the RHS of (10) is less than 2ϵ. Since ϵ is arbitrary, we have shown that $\lim_{n\to\infty} \int_{\theta\in\mathbb{R}} u_i(\mu_i, \mu_{-i}, \theta)\, d\mathbb{P}_n(\theta) = u_i^*$, so by definition, the class of nearly stationary strategies is admissible.

Necessity: We prove by contradiction that every strategy $\mu_i \in \mathcal{M}_i^0$ is nearly stationary. Since strategies for uninformed players are stationary (and nearly stationary), assume i is an informed player. Suppose toward a contradiction that $\mu_i \in \mathcal{M}_i^0$ but μ_i is not nearly stationary. By Lemma 7, there exists an attraction $\mu_i^* \in \mathcal{M}_i$ for μ_i.

Next, given the existence of an attraction μ_i^*, we establish uniqueness. Suppose there also exists an attraction $\hat{\mu}_i \in \mathcal{M}_i$ for μ_i with $\hat{\mu}_i \neq \mu_i^*$. We show that there exists a profile of stationary strategies of the rivals such that if player i plays μ_i, there is not a well-defined expected payoff in the limit. By irreducibility of payoffs (Assumption 3), there exists a state $\theta_0 \in \mathbb{R}$ and a profile μ_{-i} of stationary rival strategies such that $u(\mu_i^*, \mu_{-i}, \theta_0) \neq u(\hat{\mu}_i, \mu_{-i}, \theta_0)$. We show that this contradicts admissibility.

Note that by Lemma 6, for any $\epsilon > 0$, there exists $\eta > 0$ such that if $d_{P,i}(\mu_i^\theta, \mu_i^*) < \eta$, then

$$
\epsilon > |u_i(\mu_i^\theta, \mu_{-i}, \theta_0) - u_i(\mu_i^*, \mu_{-i}, \theta_0)|. \qquad (11)
$$

An analogous statement holds for $\hat{\mu}_i$, so let us redefine η so that for $\tilde{\mu}_i \in \{\mu_i^*, \hat{\mu}_i\}$, $d_{P,i}(\mu_i^\theta, \tilde{\mu}_i) < \eta$ implies $|u_i(\mu_i^\theta, \mu_{-i}, \theta_0) - u_i(\tilde{\mu}_i, \mu_{-i}, \theta_0)| < \epsilon$.

Recall that by admissibility, there is some u_i^* such that $\lim_{n\to\infty} \int u_i(\mu_i, \mu_{-i}, \theta) d\mathbb{P}_n(\theta) = u_i^*$ for all diffusing sequences $(\mathbb{P}_n)$. For the contradiction, we construct two sequence of measures $(\mathbb{P}_n^1)$ and $(\mathbb{P}_n^2)$ along which the limits $\lim_{n\to\infty} \int u_i(\mu_i, \mu_{-i}, \theta) d\mathbb{P}_n^1(\theta)$ and $\lim_{n\to\infty} \int u_i(\mu_i, \mu_{-i}, \theta) d\mathbb{P}_n^2(\theta)$ differ. Let $v^* = u_i(\mu_i^*, \mu_{-i}, \theta_0)$ and $\hat{v} = u_i(\hat{\mu}_i, \mu_{-i}, \theta_0)$, and recall from above that $v^* \neq \hat{v}$. Since μ_i^* is an attraction for μ_i and $\lambda\left(\{\theta : d_{P,i}(\mu_i^{\theta_0-\theta}, \mu_i^*) < \eta\}\right) = \infty$, for each $n \in \mathbb{N}$, there exists $C_n^1 \subset \{\theta : d_{P,i}(\mu_i^{\theta_0-\theta}, \mu_i^*) < \eta\} \setminus [-n, n]$ with $\lambda(C_n^1) = 2n^2$. Define $B_n^1 := [-n, n] \cup C_n^1$, and define $\mathbb{P}_n^1(\theta) := \mathbb{1}_{B_n^1}(\theta)/\lambda(B_n^1)$. By construction, $(\mathbb{P}_n^1)_{n\in\mathbb{N}}$ is a diffusing sequence of measures, and by the assumption that μ_i is admissible, we must have $\lim_{n\to\infty} \int u_i(\mu_i, \mu_{-i}, \theta) d\mathbb{P}_n^1 = u_i^*$. Pick any $C > |v^*|$ such that C is an upper bound, over all θ, on the magnitude of $u_i(\mu_i, \mu_{-i}, \theta) = u_i(\mu_i^{\theta_0-\theta}, \mu_{-i}, \theta_0)$.[14] We have

$$
\begin{aligned}
\left| \int_{\theta\in\mathbb{R}} u_i(\mu_i, \mu_{-i}, \theta) d\mathbb{P}_n^1(\theta) - v^* \right| &= \left| \int_{\theta\in\mathbb{R}} u_i(\mu_i^{\theta_0-\theta}, \mu_{-i}, \theta_0) d\mathbb{P}_n^1(\theta) - v^* \right| \\
&\leq \int_{\theta\in\mathbb{R}} \left| u_i(\mu_i^{\theta_0-\theta}, \mu_{-i}, \theta_0) - v^* \right| d\mathbb{P}_n^1(\theta) \\
&= \int_{\theta\in B_n^1} \left| u_i(\mu_i^{\theta_0-\theta}, \mu_{-i}, \theta_0) - v^* \right| d\mathbb{P}_n^1(\theta) \\
&= \int_{\theta\in C_n^1} \left| u_i(\mu_i^{\theta_0-\theta}, \mu_{-i}, \theta_0) - v^* \right| d\mathbb{P}_n^1(\theta) \\
&\quad + \int_{\theta\in[-n,n]} \left| u_i(\mu_i^{\theta_0-\theta}, \mu_{-i}, \theta_0) - v^* \right| d\mathbb{P}_n^1(\theta) \\
&\leq \frac{\epsilon \cdot 2n^2 + 2C \cdot 2n}{2n^2 + 2n} \to \epsilon,
\end{aligned}
$$

where the final inequality uses (i) that, by the earlier construction, $\theta \in C_n^1$ implies $d_{P,i}(\mu_i^{\theta_0-\theta}, \mu_i^*) < \eta$ which implies $|u_i(\mu_i^{\theta_0-\theta}, \mu_{-i}, \theta_0) - v^*| < \epsilon$ and (ii) that for all θ, $C > \max\{|v^*|, u_i(\mu_i^{\theta_0-\theta}, \mu_{-i}, \theta_0)\}$. It follows that $u_i^* \in [v^* - \epsilon, v^* + \epsilon]$.

Likewise, $\hat{\mu}_i$ is an attraction, so $\lambda\left(\{\theta : d_{P,i}(\mu_i^{\theta_0-\theta}, \hat{\mu}_i) < \eta\}\right) = \infty$, and choose $C_n^2 \subset \{\theta : d_{P,i}(\mu_i^{\theta_0-\theta}, \hat{\mu}_i) < \eta\} \setminus [-n, n]$ with $\lambda(C_n^2) = 2n^2$, $B_n^2 := [-n, n] \cup C_n^2$, and $\mathbb{P}_n^2(\theta) := \mathbb{1}_{B_n^2}(\theta)/\lambda(B_n^2)$, such that $|\int u_i(\mu_i, \mu_{-i}, \theta) d\mathbb{P}_n^2 - \hat{v}| \leq \frac{\epsilon \cdot 2n^2 + 2C \cdot 2n}{2n^2 + 2n} \to \epsilon$, and thus $u_i^* \in [\hat{v} - \epsilon, \hat{v} + \epsilon]$. Since η is arbitrary, we choose $\epsilon < \frac{|v^* - \hat{v}|}{2}$ and obtain a contradiction of the fact that $v^* \neq \hat{v}$. Hence we conclude that there is a unique attraction μ_i^*.

We now prove that μ_i^* is a limit strategy for μ_i in the sense of Definition 3. We derive a contradiction by showing that otherwise, the uniqueness result above would be violated. Suppose by way of contradiction that μ_i^* is not a limit strategy of μ_i, and hence by Lemma 2, for some $\eta > 0$, $\lambda(\Theta_{\geq \eta}) = \infty$ where $\Theta_{\geq \eta} := \{\theta : d_{P,i}(\mu_i^\theta, \mu_i^*) \geq \eta\}$. Let $\mathcal{Q} := \{\tilde{\mu}_i \in \mathcal{M}_i : \exists \theta \in \Theta_{\geq \eta} \text{ s.t. } \tilde{\mu}_i = \mu_i^\theta\}$. By (the more general result shown in the proof of) Lemma 7, there exists an attraction $\hat{\mu}_i \in \mathcal{Q}$ for μ_i. By construction, $\hat{\mu}_i \neq \mu_i^*$. This contradicts the uniqueness of the attraction μ_i^* as argued previously, so μ_i^* must be a limit strategy for μ_i, as desired.

Finally, μ_i^* is stationary by Lemma 3, and μ_i is nearly stationary. Payoff equivalence between μ_i and μ_i^* follows from a straightforward argument similar to the one given for the sufficiency part of the proof.

Proof (Proof of Theorem 2.) For sufficiency, consider any profile of nearly stationary strategies $\mu \in \mathcal{M}$. Lemma 3 applies, and hence there exists a unique profile of stationary strategies $(\mu_1^*, \mu_2^*, \ldots, \mu_N^*)$ such that μ_i^* is a limit strategy of μ_i for each $i \in \mathcal{I}$. Since $\mathcal{M}$ is closed, $\mu^* \in \mathcal{M}$. By continuity of each u_i in strategies, as $\theta \to \pm\infty$, we have $u_i(\mu, \theta) = u_i(\mu^{-\theta}, 0) \to u_i^* := u_i(\mu^*, 0)$. Now for any $\epsilon > 0$, choose $\underline{\theta}, \overline{\theta} \in \mathbb{R}$ with $\underline{\theta} < \overline{\theta}$ such that $\theta \notin [\underline{\theta}, \overline{\theta}] \implies |u_i(\mu^{-\theta}, 0) - u_i^*| < \epsilon$ for all $i \in \mathcal{I}$. For any diffusing sequence, $\mathbb{P}_n([\underline{\theta}, \overline{\theta}]) \to 0$. Since $u_i(\cdot, 0)$ is bounded in magnitude by some M, the rest of the argument for sufficiency in the proof of Theorem 1 applies.

[14] Here we have used that $\mu_{-i}^{\theta_0-\theta} = \mu_{-i}$ by the stationarity of μ_{-i}.

For necessity, suppose that there exists $i \in \mathcal{I}$ and $\mu_i \in \mathcal{M}_i^0$ such that μ_i is not nearly stationary. Note that by Assumption 4, an extension of the arguments used in the proof of Lemma 1 shows that each space of measures satisfying Definition 7 is compact, and since each $\mathcal{M}_i$ is a closed subspace, each $\mathcal{M}_i$ is also compact. Hence, for each i there exists an attraction $\mu_i^* \in \mathcal{M}_i$ in the sense of Lemma 7. If μ_i is not nearly stationary, then by definition μ_i^* is not a limit strategy. Hence, there exists $\eta > 0$ such that the set $\Theta_{\geq\eta} := \{\theta \in \mathbb{R} : d_i(\mu_i^\theta, \mu_i^*) \geq \eta\}$ satisfies $\lambda(\Theta_{\geq\eta}) = \infty$. Since the set $\mathcal{M}_{i,\geq\eta} := \{\tilde{\mu}_i \in \mathcal{M}_i : d_i(\tilde{\mu}_i, \mu_i^*) \geq \eta\}$ is a closed subspace of $\mathcal{M}_i$, it is compact, and there exists an attraction $\hat{\mu}_i \in \mathcal{M}_{i,\geq\eta}$ for μ_i. Since $\mu_i^* \notin \mathcal{M}_{i,\geq\eta}$ by construction, $\mu_i^* \neq \hat{\mu}_i$. Hence, there exists $\theta_0 \in \mathbb{R}$ and a profile of stationary strategies $\mu_{-i} \in \mathcal{M}_{-i}^0$ for the remaining players such that $v^* := u_i(\mu_i^*, \mu_{-i}, \theta_0) \neq \hat{v} := u_i(\hat{\mu}_i, \mu_{-i}, \theta_0)$. Now choose any $\epsilon > 0$ such that $\epsilon < |v^* - \hat{v}|/2$. Since payoffs are continuous in strategies, and the space $\mathcal{M}_i$ is compact, this continuity is uniform, so we can choose $\eta' > 0$ sufficiently small that $\tilde{\mu}_i \in \{\mu_i^*, \hat{\mu}_i\}$, $d_{P,i}(\mu_i^\theta, \tilde{\mu}_i) < \eta'$ implies $|u_i(\mu_i^\theta, \mu_{-i}, \theta_0) - u_i(\tilde{\mu}_i, \mu_{-i}, \theta_0)| < \epsilon$. Then, by the same construction as in the proof of Theorem 1, there exist two diffusing sequences $(\mathbb{P}_n^1)$ and $(\mathbb{P}_n^2)$ along which ex ante payoffs for player i from playing μ_i against μ_{-i} converge to limits in $(v^* - \epsilon, v^* + \epsilon)$ and $(\hat{v} - \epsilon, \hat{v} + \epsilon)$, respectively. As these are disjoint intervals, the limits are distinct, and hence μ_i is not part of an admissible class.

Proof (Proof of Proposition 1.) Stationarity of signals and payoffs are immediate from the definition of the game, as is compactness (Assumption 4), and finiteness of payoffs follows from the quadratic payoff structure with normally distributed signals. For concreteness, let the marginal G in the senders' strategies be simply $N(0, \sigma^2)$, the distribution of signals conditional on $\theta = 0$. Suppose the class $\mathcal{M}$ satisfies parts (i)-(iv) of Assumption 5; we establish irreducibility. Irreducibility for the principal is straightforward, since for two distinct principal strategies, one can marginally adjust one or both of the sender strategies and change the principal's payoffs under those two strategies by different amounts. We show irreducibility for sender 1, which by symmetry implies irreducibility for sender 2. Suppose the principal always chooses the lower offer: given offers a_1 and a_2, she chooses $C(a_1, a_2) = \arg\min\{a_1, a_2\}$. Consider two distinct strategies of expert 1, μ_1 and $\hat{\mu}_1$. We consider only $b_1 = 0$; similar arguments apply for any b_1. We show that there exists a constant markup strategy κ_m for player 2 such that $u_1(\mu_1, \kappa_m, C, 0) \neq u_1(\hat{\mu}_1, \kappa_m, C, 0)$. By the definition of payoffs and signals in the model, $u_1(\mu_1, \kappa_m, C, 0) = -\int_{\mathbb{R}^2} (\min\{a_1, s_2 + m\})^2 (Q \otimes \Phi)(d(a_1, s_2))$, where Q is the CDF over player 1's action induced by the strategy μ_1, defined by $Q(x) = \mu_1(\{a_1 : a_1 \leq x\})$, and where Φ is the CDF of $N(0, \sigma^2)$. Define $\hat{Q}(x)$ for $\hat{\mu}_1$ likewise. Next, we show that we can differentiate w.r.t. m under the integral. Write $\mathbf{a} := (a_1, s_2)$, $g(\mathbf{a}, m) := (\min\{a_1, s_2 + m\})^2$, and $\nu = Q \otimes \Phi$. Note that for all $\mathbf{a} \in \mathbb{R}^2$, $g(\mathbf{a}, m)$ is absolutely continuous in m on bounded intervals. We have

$$a(m) := \int_{\mathbb{R}^2} g(\mathbf{a}, m)\nu(d\mathbf{a}) \tag{12}$$

$$= \int_{\mathbb{R}^2} \left[g(\mathbf{a}, m_0) + \int_{m_0}^{m} g_m(\mathbf{a}, z) dz \right] \nu(d\mathbf{a}) \tag{13}$$

$$= \int_{\mathbb{R}^2} g(\mathbf{a}, m_0)\nu(d\mathbf{a}) + \int_{m_0}^{m} \int_{\mathbb{R}^2} g_m(\mathbf{a}, z)\nu(d\mathbf{a}) dz \tag{14}$$

$$\implies a'(m) := \int_{\mathbb{R}^2} g_m(\mathbf{a}, m)\nu(d\mathbf{a}) \qquad \text{a.e. } m, \tag{15}$$

where $m_0 < m$ can be chosen arbitrarily. To obtain (13) we have used the fact that absolutely continuous functions are the integral of their derivatives.[15] Fubini's Theorem is used to obtain (14). Differentiability-a.e. of the integral yields (15).[16] By the same arguments, we obtain $a''(m) = \int g_{mm}(\mathbf{a}, m)\nu(d\mathbf{a}) = 2\int_{\{\mathbf{a}: a_1 > s_2 + m\}} \nu(d\mathbf{a})$ a.e. m. Applying this to $\nu = Q \otimes \Phi$ and $\hat{\nu} = \hat{Q} \otimes \Phi$, it follows that if $u_1(\mu_1, \kappa_m, C, 0) = u_1(\hat{\mu}_1, \kappa_m, C, 0)$ for all $m \in \mathbb{R}$, then $\int_{\{\mathbf{a}: a_1 > s_2 + m\}} (Q \otimes \Phi)(d(a_1, s_2)) = \int_{\{\mathbf{a}: a_1 > s_2 + m\}} (\hat{Q} \otimes \Phi)(d(a_1, s_2))$ a.e. m. Rearranging, we have

$$\int_{S_2} (1 - Q(s_2 + m)) f_2(s_2) ds_2 = \int_{S_2} (1 - \hat{Q}(s_2 + m)) f_2(s_2) ds_2$$

$$\implies 0 = \int_{S_2} (Q(m - s_2) - \hat{Q}(m - s_2)) f_2(s_2) ds_2, \tag{16}$$

by the evenness of the normal distribution, where f_2 is the PDF of $N(0, \sigma^2)$, the noise distribution for player 2. Letting $K := Q - \hat{Q}$, (16) is a convolution equation:

$$[K * f](m) = 0 \qquad \text{a.e. } m.$$

Let $\mathcal{F}$ denote the normalized Fourier transform, $\mathcal{F}(g)(z) := \int_{-\infty}^{\infty} e^{-2\pi i z x} g(x) dx$. Both K and f are Lebesgue integrable functions, and thus by the convolution theorem,[17] $\mathcal{F}(K * f) = \mathcal{F}(K) \cdot \mathcal{F}(f)$. But since $K * f \equiv 0$, $\mathcal{F}(K * f) \equiv 0$. Since $\mathcal{F}(f)(z) = e^{-2(\pi\sigma z)^2} > 0$ for all z, we must have $\mathcal{F}(K) \equiv 0$. Applying the inverse Fourier transform, we have $K = 0$ almost everywhere. This contradicts the assumption $Q \neq \hat{Q}$, so it must be that $u_1(\mu_1, \kappa_m, C, 0) \neq u_1(\hat{\mu}_1, \kappa_m, C, 0)$ for some m.

References

Ambrus A, Baranovskyi V, Kolb A (2019) A delegation-based theory of expertise. Unpublished working paper

Angeletos GM, Lorenzoni G, Pavan A (2010) Beauty contests and irrational exuberance: A neoclassical approach. Tech. rep., National Bureau of Economic Research

[15] See Royden and Fitzpatrick (1988, Corollary 5.15).

[16] See Royden and Fitzpatrick (1988, Theorem 5.10).

[17] See Reed and Simon (1980, Theorem IX.3b).

Billingsley P (2009) Convergence of probability measures, vol 493. John Wiley & Sons
Dale DJ, Morgan J (2015) A comment on “The welfare effects of public information” by Stephen Morris and Hyun Song Shin. mimeo, University of California, Berkeley
Friedman D (1991) A simple testable model of double auction markets. Journal of Economic Behavior and Organization 15:47–70
Hellwig C, Veldkamp L (2009) Knowing what others know: Coordination motives in information acquisition. The Review of Economic Studies 76(1):223–251
Jeffreys H (1946) An invariant form for the prior probability in estimation problems. Proceedings of the Royal Society of London Series A Mathematical and Physical Sciences 186(1007):453–461
Jeffreys H (1961) The theory of probability. Oxford University Press
Kass RE, Wasserman L (1996) The selection of prior distributions by formal rules. Journal of the American Statistical Association 91(435):1343–1370
Klemperer P (1999) Auction theory: a guide to the literature. Journal of Economic Surveys 13:227–286
Laplace PS (1951) A philosophical essay on probabilities, 1819. English translation, Dover 6
Milgrom P, Weber R (1985) Distributional strategies for games with incomplete information. Mathematics of Operations Research 10:619–632
Morris S, Shin H (2002) The social value of public information. The American Economic Review 92:1521–1534
Morris S, Shin H (2003) Global games: Theory and applications. Econometrics Society Monographs
Myatt D, Wallace C (2014) Central bank communication design in a Lucas-Phelps economy. Journal of Monetary Economics 63:64–79
Ok EA (2007) Real analysis with economic applications, vol 10. Princeton University Press
Reed M, Simon B (1980) Methods of modern mathematical physics: Functional analysis, vol 1. Gulf Professional Publishing
Royden H, Fitzpatrick P (1988) Real analysis, vol 4. Prentice Hall New York
Yang R, Berger JO (1998) A catalog of noninformative priors. Tech. rep.