ANNUAL REVIEWS

Annual Review of Economics

Large Games: Robustness and Stability

Ronen Gradwohl[1] and Ehud Kalai[2]

[1]Department of Economics and Business Administration, Ariel University, Ariel 40700, Israel; email: roneng@ariel.ac.il

[2]Kellogg School of Management, Northwestern University, Evanston, Illinois 60208, USA; email: kalai@kellogg.northwestern.edu

Annu. Rev. Econ. 2021. 13:3.1–3.18

The *Annual Review of Economics* is online at economics.annualreviews.org

https://doi.org/10.1146/annurev-economics-072720-042303

Copyright © 2021 by Annual Reviews.
All rights reserved

JEL codes: C72, D80

Keywords

large games, Nash equilibrium, ex-post Nash, structural robustness, fault tolerance, viability

Abstract

This review focuses on properties related to the robustness and stability of Nash equilibria in games with a large number of players. Somewhat surprisingly, these equilibria become substantially more robust and stable as the number of players increases. We illustrate the relevant phenomena through a binary-action game with strategic substitutes, framed as a game of social isolation in a pandemic environment.

1. INTRODUCTION

The vast interconnectedness of present-day economic, social, and political activities often magnifies the degrees of uncertainty and complexity that are inherent in human interaction. To study such interactions—in order to understand the underlying strategic forces, predict outcomes, and recommend policy interventions—one may turn to game theory. However, a common critique of game theory is that, while it has excelled at providing precise analyses of small, fully specified, and idealized interactions, many real-life conditions are rarely so clean. It is thus not clear whether the predictions it produces are robust and stable. Instead, they are often seen as fragile and as unable to withstand departures from full specification and rationality.

The issues arise both in normative and in descriptive analyses. From a normative perspective, when one designs mechanisms or institutions with a strategic element, robustness and stability are clear desiderata: In practice, since participants rarely interact only within the context of the particular mechanism under design, the mechanism must function well even in a broader context. From a descriptive perspective, when analyzing interactions, one would most likely expect to see robust, stable equilibria.

In this review we argue that, within the context of large games, the critique is much less salient and the robustness and stability properties much more prevalent. In particular, we focus on the most widely used concept in game theory—Nash equilibrium—and show that in large games it is sound, robust, and reasonable. Furthermore, we argue that economics needs this tool to address important economic problems. We illustrate this latter point with a timely example about social isolation.

1.1. Related Literature on Large Games

There are several somewhat overlapping branches of the literature on what might be called large strategic games.[1] The first is one that considers a continuum of players. The foundational papers in this branch focus on the existence of pure-strategy equilibria (Schmeidler 1973), but they have also found many applications to economic problems. Much of the earlier work on continuum-player games is surveyed by Khan & Sun (2002), but Khan et al. (2017), Cerreia-Vioglio et al. (2020), and the references therein offer more recent contributions.

A second branch of the literature considers a large number of players but is concerned with properties of the limit (see, e.g., Al-Najjar 2008 and the references therein). In particular, its focus is on whether the limit of these games resembles their idealization as either countably many players or as a continuum of players, and, as such, on when the idealization serves as an accurate approximation and when it does not.

This review is concerned with a third branch of the literature, which focuses on the robustness and stability of large, finite games (see, e.g., Kalai 2004 and the references therein, as well as follow-up references cited in this review). While the robustness and stability properties of such games will generally hold only when there are sufficiently many players, the requisite number of players will typically be finite and depend on primitive parameters of the game under study.

1.2. Related Literature on Mechanism Design

The literature on mechanism design is also related to the current review, both in terms of settings with many participants and in the analysis of robustness. The former case includes, for example,

[1]We do not discuss the older literature on large cooperative games (see Aumann & Shapley 1974). There, in parallel to the results reported in this review, theoretical predictions also become more robust as the number of players increases.

the work of Azevedo & Budish (2019) on incentives in large mechanisms and of Kearns et al. (2014) on recommender mechanisms. Additionally, researchers have examined specific mechanisms with many participants, including the work of Cole & Tao (2016) on market games, of Lee (2016) and Che & Tercieux (2019) on matching markets, and of Bodoh-Creed (2013) on auctions.

A different, emerging branch of the literature analyzes issues of robustness in the context of mechanism design and contract theory (see Carroll 2019 for a survey). One main difference between that work and the current review is that a major focus of the former is the design and analysis of mechanisms that are robust to beliefs, and especially to non-Bayesian uncertainty. In contrast, we are concerned here with the inherent robustness of equilibria in large games.

2. MAIN ILLUSTRATIVE EXAMPLE: SOCIAL ISOLATION

We illustrate many of our points through a binary-action game with strategic substitutes. To help keep track of the ideas, we interpret the game as one of social isolation, motivated by the recent social and policy debates surrounding the COVID-19 pandemic.

Although the behavioral and strategic aspects of disease transmission and policy responses have been studied extensively in economics,[2] the recent COVID-19 pandemic has intensified this research. Much of this work takes as its baseline the dynamic epidemiological SIR model of Kermack & McKendrick (1927), which models the change in the spread of an infection within a population over time. Researchers have extended this model to incorporate strategic behavior by the populace-i.e., –whether people decide to stay home or go out and their level of social activity—and examined various policy interventions, such as full or partial targeted lockdowns.[3] These economic models emphasize the importance of strategic behavior by individuals in driving both the spread and the containment of infection, as well as the effects of various policies.

The social isolation (henceforth, SI) games we analyze highlight this strategic behavior. Given the fundamentals of the environment, such as the costs and benefits of isolation and of social activity, will individuals tend to stay home, in social isolation, or will they go out and engage in social activities? And how will changes in the environment, brought about by government policies or changing fundamentals, affect their behavior? Our focus will be on the robustness and stability of these predictions.

In addition to illustrating the robustness and stability of large games, the examples also draw attention to the importance of these properties. In the disease transmission models described above, behavioral predictions are typically used as inputs. If the predictions are fraught with uncertainty, then the conclusions of the epidemiological models are called into question. And the behavioral predictions may well be fraught with uncertainty, as the environment is far more complex than the idealized variant: Individuals can communicate with their peers and learn from others, they can change their behavior over time, they are exposed to news and other information through social media, and much more. Thus, whether or not predictions in the idealized model remain stable and robust to such features may well be the factor that determines the viability of the epidemiological models' conclusions.

In this review we address these issues of robustness and stability. As noted above, we illustrate them through a binary-action game of strategic substitutes that we interpret as a game of social

[2] Readers are referred to, for example, Geoffard & Philipson (1996), Gaffeo (2003), Goenka & Liu (2012), Adda (2016), and Greenwood et al. (2019).

[3] Acemoglu et al. (2020), Alvarez et al. (2020), Farboodi et al. (2020), and Keppo et al. (2020) are a few representative contributions.

isolation. We consider two basic benchmark variants of the SI game: one with complete information and one with incomplete information. Both are played in an environment $\mathcal{E} = (n, h, o)$ that consists of a set N of n players, where n is a (typically large) positive integer. Each player may choose to be in one of two possible locations: at home (H) or out (O). Each player type is described by a pair of payoff functions (h_i, o_i), where $h_i \in \mathbb{R}$ represents the payoff of player $i \in N$ at home and $o_i(p) \in \mathbb{R}$ is i's payoff of being out when the proportion of the other players who are out is p. We assume that each o_i is a continuous and strictly decreasing function with $o_i(0) \geq h_i$ and $o_i(1) \leq h_i$.

2.1. Complete Information

We begin with the simplest game in the environment $\mathcal{E} = (n, h, o)$, one with complete information and symmetric utilities.

Example 1 (symmetric SI game with complete information). Consider the environment $\mathcal{E}$ in which $h_i = h \in [0, 1]$ and $o_i(p) = o(p) = 1 - p$ for all players i. That is, each player obtains utility h from staying home and utility $1 - p$ from going out, where p is the fraction of the other players who are also out.

Let $c = o^{-1}(h) = 1 - h$. Then, in the unique symmetric Nash equilibrium, each player chooses O with probability c and H with probability $1 - c$.[4]

How robust is this equilibrium, and is it still a sound prediction in nonidealized, real-world settings? For instance, if a player obtains some ex-post estimate of the number of other players going out (for example, via the news or social media), will that compel them to change their own action? This is a question about rational expectations: A player's equilibrium strategy is optimal given their expectation about the future behavior of others. But if the player were allowed to change their mind after observing others' actions, at which point some or all uncertainty would get resolved, would they still stick with their initial plan?

Relatedly, what if the player hears about some others who do not act as expected? What if they have other considerations, for example, whether or not their family or neighbors are out: Do these affect their choice of action? Are there other equilibria? How robust are they? And what happens in more complex environments in which players are not symmetric and utilities from O are not linear?

Since the predictions of the simple model are likely to be used in a policy debate and to generate policy recommendations, additional questions arise. For example, how close is the actual number of people who are out to the expected number of people out in the idealized setting? To what degree would different policies that change the costs and benefits of staying home (such as support for telecommuting or enhanced food delivery services) or going out (such as fines, restrictions, or limited transportation) affect these numbers? And how much uncertainty surrounds the effects of such policies? A policy that leads to a definite but minor decrease in the number of people out is very different from one that leads to the possibility of a major decrease but also to a significant chance of an increase, even if the expected change is the same for both.

In order to be compelling, predictions in the game should answer these questions, and we will study the conditions under which they do. Before that, however, we turn to an incomplete-information version of this game, in which the issues above are further exacerbated.

[4]To see this, fix a player i and observe that the expected fraction of other players out is $c = 1 - h$. Hence, i's expected utility from O is $o(1 - h) = h$, which is equal to their utility from H.

2.2. Incomplete Information

Example 2 (symmetric SI game with incomplete information). Consider the environment $\mathcal{E} = (n, h, o)$, in which each h_i is the type of player i. These types are drawn in an independent and identically distributed (i.i.d.) manner from a continuous, strictly increasing cumulative function L that has full support on the interval [0, 1]. Also, fix $o_i(p) = 1 - p$ for every player i.

Define the cutoff type to be the unique C such that $C = \Pr(h_i \geq C)$. Consider a symmetric cutoff strategy $\widehat{C}$ in which a player chooses H if $h_i \geq C$ and O if $h_i < C$.

Claim 1. The ex-ante symmetric profile in which every player uses $\widehat{C}$ is a Nash equilibrium.

Proof. At the profile above, the expected payoff of a player i who chooses O is $o_i(p) = 1 - E(p) = 1 - \Pr(h_j < C) = C$. So H is preferred to O if and only if $h_i \geq C$, which is precisely the player's strategy under $\widehat{C}$. □

Again, we should ask: How robust is this equilibrium, and how sound is it in nonidealized settings? The questions relating to the complete information game of Example 1 are relevant here as well. But now there are additional issues. For instance, if a player obtains some information about the types of other players (for example, via conversations with friends and family, perhaps through social media), will that compel them to change their own action?

From a policy perspective, to what extent will providing relevant information to individuals, in an attempt to change their types, affect the number of people who go out? And again, how much uncertainty is there surrounding this effect?

2.3. Asymmetric Social Isolation Games

We will use the symmetric SI games of Examples 1 and 2 to illustrate the ideas in the review. Sometimes it will be helpful to refer to more general variants of these games, and so we formally describe two of them here.

The complete-information SI game of Example 1 is symmetric, with linear utilities to players who are out. More generally, we consider an environment in which each player i has their own h_i, and where o_i is an arbitrary decreasing function of the fraction p of others who are out.

Example 3 (asymmetric SI game with complete information). Consider the environment $\mathcal{E}$, where each player i's utilities are $h_i \in \mathbb{R}$ and $o_i : [0, 1] \mapsto \mathbb{R}$. That is, each player i's utilities h_i from staying home and $o_i(p)$ from going out may be different from those of other players.

We also consider a more general version of the incomplete-information SI game of Example 2. Such a variant incorporates asymmetric distributions over types as well as different benefits and costs to being out.

Example 4 (asymmetric SI game with incomplete information). Consider the environment $\mathcal{E} = (n, h, o)$, in which each player i's type is (h_i, b_i, c_i). Types are independent across players but may be drawn from nonidentical distributions. Player i's utility from staying home is h_i and from being out is $o_i(p) = b_i - p \cdot c_i$, where b_i can be interpreted as the benefit and c_i the cost of being out.

2.4. Outline

Consider the complete-information SI game of Example 1 and observe that, when there are only two players, this is a simple coordination game. In this case, the equilibrium predictions are perhaps fragile in the sense discussed above; for example, the simultaneous-move assumption may be questionable, since a player in such a game would prefer to wait and observe the choice of their opponent before choosing their own action. At the other end of the spectrum, one might consider a variant in which the population of players is a continuum. In this case, adopting an assumption of price-taking behavior—an approach common in economic theory—might make sense. However, when one designs policies for social isolation, the relevant communities might be smaller and hence not adequately modeled as a continuum. Additionally, the assumption in price-taking models that individuals exert zero influence on others is not always valid.

The research on large games that is surveyed in this article bridges the two ends of the spectrum. On the one hand, it allows us to retain the structure and simplicity of price-taking models, but without imposing price-taking behavior as an assumption nor assuming a continuum of players. On the other hand, the precise equilibrium analysis it facilitates is robust and stable—when the number of players is large.

We proceed as follows. In Section 3 we describe the properties of games that render them large and the conditions under which their equilibria are robust and stable. We show that these conditions hold in the simple SI games above but also in rather broad generalizations of these baselines. In Section 4 we then describe the stability and robustness properties inherent in large games by illustrating them in variants of the SI game. This is accompanied by formalization of these properties, a description of the relationships between them, and an examination of the necessary conditions on games under which they hold. Finally, we conclude in Section 5 with a discussion of directions for further research.

3. CLASSES OF GAMES

In trying to analyze strategic situations, predict their outcomes, or design institutions for strategic participants, one may hope that the sheer size of these interactions renders their equilibria robust and stable (in the sense to be described in Section 4). Of course, a game is not large simply by virtue of the number of players—one can always take a small game and add dummy players. To address this issue, one might consider different classes of games for which the robustness and stability properties hold.

We begin by discussing a class of games that has been studied extensively in the literature as representing largeness, namely, anonymous games. We then turn to another class, which features the property of low strategic sensitivity that has been identified as crucial to guarantees of robustness and stability. We show that the simple SI games of Examples 1 and 2 belong to both classes, but various extensions and generalizations might not.

3.1. Anonymous Games

A game is anonymous if a player's utility depends on their own type and action as well as on the realized empirical distribution of other players' type-action pairs. More formally, for any player i, any pair of other players j and k, any type profile t, and any action profile a, the utility of player i is $u_i(t, a) = u_i(t', a')$, where t' and a' are identical to t and a except that the type-action pairs of players j and k are swapped. We note that Kalai (2004) calls this condition semianonymous to emphasize that the utility of player i may depend on i's specific type and action, and that i's utility remains unchanged when the types and actions of players j and k who are different from i are swapped.

The symmetric SI games of Examples 1 and 2, as well as the asymmetric variants of Examples 3 and 4, are anonymous games, since each player's utility depends only on the fraction of others who are out and not on their identities. For an example of a game that is not anonymous, one might consider network game variants (Galeotti et al. 2010) of these examples, in which players' utilities from being out are determined not by the fraction of all others, but only by a fraction of neighbors in some social network who are out.

A useful generalization of anonymous games is that of aggregative games, which capture many applications in industrial organization, macroeconomics, political economy, and public economics (see, e.g., Jensen 2010, Acemoglu & Jensen 2013). In aggregative games, a player's payoff depends on their own action and some function that aggregates others' actions. Anonymous games are the special case in which the (multidimensional) aggregator is the fraction of players who play each action. Because of their importance, many properties of anonymous and aggregative games have been examined, such as the existence of pure equilibria (e.g., Jensen 2010), comparative statics (e.g., Acemoglu & Jensen 2013), and equilibrium characterization and computation (e.g., Blonski 1999, Daskalakis & Papadimitriou 2015).

3.2. Games with Low Strategic Sensitivity

The strategic sensitivity of a game is a measure of the influence that players' types and actions have on other players' utilities. A game features low strategic sensitivity if players do not have an outsized influence on others' utility functions. The most extreme example is the case in which we have n decision problems, i.e., each player's payoff depends only on their own type and action. A less extreme case is illustrated by Example 1, which also has low strategic sensitivity: If player i chooses H, others' actions have no effect on their payoff. Moreover, if player i chooses O, then a change in the action of some other player leads to a change in player i's utility from $o_i(p)$ to $o_i(p')$, where $p' = p + 1/(n-1)$ or $p' = p - 1/(n-1)$. In Example 1 we assumed that $o_i(p) = 1 - p$, so the change in utility is a gain or loss of at most $1/(n-1)$.

There are various ways to define strategic sensitivity. Here we present the variant of Azrieli & Shmaya (2013), which they call the Lipschitz constant of the game. Other variants, some slightly stronger and some slightly weaker, were utilized by, for example, Milgrom & Weber (1985), Kalai (2004), Deb & Kalai (2015), and many others, sometimes under the name of continuity or equicontinuity of payoffs.

Definition 1 (Lipschitz constant). The Lipschitz constant of a complete information game G is $\delta(G) = \max\left\{\left|u_i(a_i, a'_{-i}) - u_i(a_i, a''_{-i})\right|\right\}$, where the maximum is taken over all i, a_i, a'_{-i}, and a''_{-i}, subject to the constraint that a'_{-i} and a''_{-i} differ on exactly one coordinate.

In words, the Lipschitz constant of a game is the maximal change one player's action can cause to the utility of another player, given any action profile.

The analysis at the beginning of this subsection demonstrates that the Lipschitz constant of the game in Example 1 is $1/(n-1)$. In contrast, the asymmetric SI game of Example 3 may have a larger Lipschitz constant. For example, if there is a player i for whom $o_i(1/2) = 1/2$ and $o_i(1/2 + 1/(n-1)) = 3/4$, then the Lipschitz constant is at least 1/4. Nonetheless, if for all i's the function o_i has bounded slope, then for a sufficiently large number of players low strategic sensitivity will be satisfied.[5]

[5] More formally, suppose there is some $M \geq 0$ such that $o'_i(p) > -M$ for all i's and p's. Then, for any n and $p \in [0, 1 - 1/(n-1)]$, the difference $o_i(p) - o_i(p + 1/(n-1)) < M/(n-1)$. Thus, for any $\varepsilon > 0$, the Lipschitz constant of the game is at most ε whenever $n > M/\varepsilon + 1$.

Furthermore, a game may have small Lipschitz constant despite not being anonymous. For example, the network game variant (Galeotti et al. 2010) of Example 1 discussed in Section 3.1 above, in which each player's utility from being out is determined by the fraction of neighbors in a social network who are out, is not anonymous. Nonetheless, if all players have many neighbors in the social network, then this game has low strategic sensitivity.

The definition of a Lipschitz constant can be generalized to games of incomplete information by defining $\delta(G) = \max\left\{\left|u_i(r_i, r'_{-i}) - u_i(r_i, r''_{-i})\right|\right\}$, where r_j denotes the type-action pair of player j, and where r'_{-i} and r''_{-i} differ on exactly one player's type-action pair. For example, the incomplete information SI game of Example 2 has a Lipschitz constant of $1/(n-1)$: For any realization of types and actions, changing one player's type and/or action leads to a change in at most $1/n$ of the fraction of players out, and so it changes the utility of a player who is out by at most $1/(n-1)$ (and that of a player who is home by zero).

4. ROBUSTNESS AND STABILITY

In this section we delve into four robustness and stability properties of large games: the existence of pure-strategy equilibria, hindsight stability, structural robustness, and fault tolerance. We illustrate each property within the context of the SI games. We then formalize the properties, discuss the relationships among them, and examine the conditions on games under which they hold.

4.1. Pure Equilibria

A first desideratum of complex games is that their equilibria be simple. Due to the complexities and conceptual issues associated with mixed-strategy equilibria (see, e.g., Rubinstein 1991), one might ask if there always exists an equilibrium in pure strategies. This, in fact, is the main focus of the older literature on games with a continuum of players (Schmeidler 1973). But what if the number of players is large but not infinite?

Let us first return to the complete-information SI game of Example 1, and recall that its symmetric equilibrium is not pure. However, this game does have a pure, asymmetric equilibrium.[6] In fact, pure equilibria also always exist in the more general, asymmetric SI game of Example 3, in which players' h_i values are possibly distinct and their o_i functions are arbitrary. Readers are referred to the **Supplemental Appendix** for a proof of this fact.

The existence of pure equilibria in large games goes well beyond SI games, especially if we also consider pure approximate equilibria. Recall that a strategy profile σ in a normal-form game G is an ε-Nash equilibrium if no player can increase their utility by more than ε by deviating. Many authors have studied various necessary conditions for pure ε-Nash equilibria to exist. Papers that study games whose type and action spaces are finite include those by Rashid (1983), Wooders et al. (2006), and Cartwright & Wooders (2009). Carmona (2008) considers compact type and action spaces and shows that pure approximate equilibria exist in games that are anonymous and have low strategic sensitivity. Azrieli & Shmaya (2013) show that low strategic sensitivity suffices for the existence of such equilibria, but a stronger result holds (in the sense of not requiring strategic sensitivity to be as low) in games that also satisfy anonymity. Finally, Carmona & Podczeck (2020) study the existence of strict pure strategy Nash equilibria and show that such equilibria exist in large games that are anonymous, have low strategic sensitivity, and satisfy some additional differentiability conditions.

[6] In particular, let $\overline{C}_n = \lfloor (n-1) \cdot C \rfloor$ and consider the pure profile in which $\overline{C}_n + 1$ players choose O and the rest choose H.

In a further generalization, one might consider strategies that are simple but not quite as simple as pure ones: namely, strategies that are mixed but have small support. Additionally, one might require that the strategies be mixed uniformly over that small support (e.g., Lipton et al. 2003, Babichenko et al. 2014). For example, Babichenko et al. (2014) show that any finite game has an approximate equilibrium in which players mix uniformly over roughly $\log n + \log m$ actions, where m is the size of each player's action space.

4.2. Self-Purification, Hindsight Stability, and Rational Expectations

Let us return to the complete-information SI game of Example 1. On further examination of its symmetric equilibrium, we may observe a stronger property than the existence of a pure equilibrium. Since players randomly choose whether to go out or stay home, the expected fraction of players who choose O is c. But the law of large numbers guarantees that, for sufficiently many players, the realized fraction of players who choose O is close to c. Note also that when the fraction of other players who choose O is close to c, the utilities of a player from choosing action H and from choosing action O are very close. The implication is that in this case, the realized profile of actions forms a pure approximate Nash equilibrium. The symmetric, mixed equilibrium in the game is thus self-purifying (Kalai 2004): Nearly all realizations of the symmetric equilibrium are, in fact, pure approximate equilibria themselves.

An alternative interpretation of self-purification is that each player, after choosing their own action and observing the realized actions of the others, can gain little by revising their choice. In Example 1, the symmetric, mixed equilibrium is thus hindsight stable or ex-post Nash.

More formally, a strategy profile is (ε, ρ)-stable if the probability that the realized action profile is one in which some player can unilaterally deviate and improve their payoff by more than ε is at most ρ. In the symmetric equilibrium of Example 1, tight concentration-of-measure bounds, such as those of Chernoff or Hoeffding (see, e.g., Alon & Spencer 2004), imply that the probability that the realized fraction of players who are out is not within $[c - \delta, c + \delta]$ is roughly $e^{-\delta n}$. Thus, for a $1 - e^{-\delta n}$ fraction of realizations of the symmetric, mixed profile, each player's difference in utility between O and H is at most δ. The probability that this difference in utilities is small for all players simultaneously is at least $1 - ne^{-\delta n}$, implying that the symmetric equilibrium here is (δ, ρ)-stable, with $\rho = ne^{-\delta n}$. Note that for any fixed $\delta > 0$, the probability that any player has a profitable deviation in the realized profile decreases to zero exponentially with n.

The notion of hindsight stability can be extended to games of incomplete information. Consider the incomplete-information SI game of Example 2, and observe that the expected number of players whose type $b_i \geq C$ is C. Again, concentration bounds guarantee that the realized number of players with such a type is around C. Thus, after the game is played, a player who observes the realized types and chosen actions of the others probably has little to gain by deviating from their own equilibrium strategy.

Hindsight stability is closely related to rational expectations. Loosely, having rational expectations means that after uncertainty has been resolved, and after what happens in some interaction has been observed, it is optimal for the economic agents to proceed with the planned continuation strategy (see, e.g., Jordan & Radner 1982). But the condition that "what happens has been observed" can be specified in different ways. Hindsight stability can be viewed as rational expectations, assuming that players can revise their strategies ex post, after observing the realized types and chosen actions of the other players.

As is true with purification, this notion of hindsight stability or rational expectations holds in much more general games. Kalai (2004) studies games (*a*) that satisfy both anonymity and low strategic sensitivity, (*b*) in which types are independent, and (*c*) in which the type and action spaces are finite, and shows that all their equilibria are (ε, ρ)-stable, where $\varepsilon > 0$ is arbitrarily small and ρ decreases to zero exponentially with n. Carmona & Podczeck (2012) generalize this result to games with infinite type and action spaces. Deb & Kalai (2015) dispense with anonymity and show that the hindsight stability result holds in games that have low strategic sensitivity and in which the type and action spaces are finite-dimensional Euclidean spaces.

We can apply the results of Carmona & Podczeck (2012) or Deb & Kalai (2015) to the asymmetric incomplete-information game of Example 4, in which each player's parameters are chosen independently of, but not necessarily identically to, those of other players. The results imply that, for a sufficiently large number of players, all equilibria of this asymmetric SI game are hindsight stable.

Despite the generality of the results, hindsight stability generally holds only in games with low strategic sensitivity. To address this concern, Gradwohl & Reingold (2010) study games in which low strategic sensitivity may fail and propose a weakening of hindsight stability called partial hindsight stability: Rather than allowing players to revise their strategies after observing the realized types and chosen actions of all others, they consider the possible benefit to a player from deviating after observing the types and actions of a randomly chosen subset of others. They show that, in all equilibria of all large games with finite action spaces and independent types, if the size of the observed subset of players is not too large, then with high probability the possible benefit from deviation is small.

Although this result does not restrict players' utility functions, it does require types to be independent. This is a main limitation also of the results of Kalai (2004), Carmona & Podczeck (2012), and Deb & Kalai (2015) described above. In Section 5 we discuss some recent progress along this front.

We conclude this section with a final observation about the policy implications of hindsight stability. In the SI games, for example, policies might change players' incentives in order to affect the equilibrium fraction of players who are out. However, policies are differentiated not only by the expected proportion of people they incentivize to go out but also by the stability of this proportion. So, for example, a policy that leads to 1/2 of the population going out is very different from one that leads to an equal chance of 1/4 or 3/4 of the population going out.

Hindsight stability can help here. Consider again the complete-information SI game of Example 1. Hindsight stability implies that in any equilibrium, with high probability the fraction of people who are out is very close to $c = 1 - h$. Now consider a policy that, for example, increases the benefit of staying home to, say, $h' = h + \alpha$. Then, hindsight stability implies that in any equilibrium, the fraction of people who are out is very close to $1 - h - \alpha$. A similar conclusion holds in the incomplete-information SI game of Example 2 as well as in the asymmetric SI games of Section 2.3.[7]

4.3. Structural Robustness

Consider the SI game, but enriched this time with additional options. For example, before deciding whether to go out or not, a player may read political forecasts and communicate with others by phone and on Facebook. If a player decides to go out, after stepping out and seeing the traffic in the street, the player may change their mind and step back into their home.

[7] For example, in the game of Example 3, in any equilibrium the fraction of players who are out is very likely close to c_q, where q is the quantile player described in the **Supplemental Appendix**. Different policies could lead to changes in the identity of that player—for example, to q'—with the result that the equilibrium fraction of players out would very likely be close to $c_{q'}$.

The notion of structural robustness captures these possible changes in the order of play, the leakage of information and communication between players about realized types and actions, and the addition of other actions, such as delegation of choices to opponents and commitments to choose certain actions. Specifically, an equilibrium profile is structurally robust if it remains an (approximate) equilibrium despite such changes.

Alternatively, structural robustness may be viewed as a stronger version of rational expectations: However the game is enriched with changes, the players' original strategies remain optimal, even when conditioned on the individual information and actions available at all the information sets of the enriched game.

More formally, a Nash equilibrium profile σ in a normal-form game G is structurally robust if an appropriate adaptation σ^A of σ remains an equilibrium (or an approximate equilibrium) in every alteration G^A of G. A structural alteration of G is any extensive-form game G^A that satisfies the following conditions (adapted from Kalai 2008):

1. Original players are included: The players in G^A are a superset of the players in G.
2. Type structure is unaltered: The distribution over the types of players in G^A, restricted to those players also in G, is the same as the distribution over those types in G.
3. Playing G^A means playing G: For every payoff node z in G^A there is an associated pure profile $a(z) = (a_1(z), \ldots, a_n(z))$ in G, such that the payoffs to the G players at z is equal to their payoffs in G under $(a_1(z), \ldots, a_n(z))$.
4. Original strategies are preserved: Every pure strategy a_i of player i in G has at least one adaptation to G^A, namely, a strategy a_i^A that guarantees that, regardless of the strategies of other players, play ends at a node z for which $a_i(z) = a_i$.

The appropriate adaptation σ^A of σ to the game G^A is the profile of mixed strategy in which each player i plays a_i^A in G^A with the same probability with which they play a_i in G under σ.

As is true for approximate notions of equilibrium, it is helpful to also consider approximate versions of structural robustness. In particular, an equilibrium σ in G is (ε, ρ)-structurally robust if, for every player i and in every alteration G^A, the probability that i reaches an information set of G^A in which they can deviate and improve their payoff by more than ε is at most ρ. This probability is taken over (*a*) the distributions of types of all players and (*b*) the strategies of all players other than i. Of course, for this to be a relevant notion we want both ε and ρ to be as small as possible. Kalai (2004) proves the following structural-robustness theorem: If a game is anonymous and has low strategic sensitivity, finite type and action spaces, and independent types, then its equilibria are (ε, ρ)-structurally robust, with $\varepsilon > 0$ arbitrarily small and ρ tending to zero exponentially as the number of players grows large.

The properties of purification and hindsight stability discussed in the previous sections are special cases of structural robustness. To see this, consider the alteration G^A of G in which players play the game in two stages. In the first stage, they each announce a type and an action (t_i, a_i^1) from G; then, after being fully informed of the outcomes of the first stage, each chooses a final action a_i^2 from G. Consider any equilibrium σ of G with an adaptation σ^A in which in the first stage players reveal their true types and use their σ_i strategy to choose their a_i^1 actions; then, in the second stage, the players do not revise their action but rather choose $a_i^2 = a_i^1$. Structural robustness implies that if σ is an (approximate) equilibrium of the normal-form game, then each player has an incentive to play σ^A in the altered game, which implies (approximate) hindsight stability.

The relationship between structural robustness and hindsight stability is actually bidirectional. In particular, Kalai (2004) shows that the converse of the above also holds, namely, that asymptotic hindsight stability implies asymptotic structural robustness. Thus, subsequent results about hindsight stability also apply. Applied to Examples 1–4, this conclusion implies that all the equilibria of these games are not only hindsight stable but also structurally robust.

4.4. Fault Tolerance

Consider again our player in the SI game, contemplating whether or not to go out. But now suppose the player hears that some of their neighbors, against all expectations and contrary to their own best interests, insist on going out. Given this unexpected, nonequilibrium play, should our player stick to their initial decision?

These questions relate to notions of fault tolerance from cryptography and distributed computing, which study the design of mechanisms that are robust against unexpected, faulty behavior by a subset of participants. In our setting, we are interested in the fault tolerance of equilibria: Is a player's equilibrium strategy still optimal even if a subset of other players deviate arbitrarily?

One version of fault tolerance is well known in economics: Specifically, a dominant-strategy equilibrium is a profile in which each player's strategy is optimal regardless of the actions of others. Hence, it is fault-tolerant to deviations by all other players. Dominant-strategy equilibria are often desirable in economic design and analysis, and it is widely recognized that such equilibria are more realistic than Nash equilibria. The problem, of course, is that often they do not exist. Fault tolerance can then be viewed as spanning the spectrum between dominant-strategy equilibria and Nash equilibria. They are optimal regardless of the actions of a subset of the other players, but when the remaining players do play according to their Nash equilibrium strategies. The level of fault tolerance is then equal to the maximal number of players who may play arbitrary actions; specifically, a t-fault-tolerant equilibrium is one in which the equilibrium strategy of any player remains optimal when at most t others deviate arbitrarily. A dominant-strategy equilibrium is thus $(n-1)$-fault-tolerant. Furthermore, an approximate version of fault tolerance is also often useful: A (t, ε)-fault-tolerant equilibrium is one in which no player can gain more than ε by deviating, regardless of the actions of any t other players, as long as the remaining $n - t - 1$ players play their equilibrium strategies.

Although notions of fault tolerance for strategic settings originated in mechanism design frameworks (see Eliaz 2002, Abraham et al. 2006), the equilibria of many large games turn out to be naturally fault-tolerant. In particular, Gradwohl & Reingold (2014) show that the equilibria of large games are (t, ε)-fault-tolerant, where $\varepsilon > 0$ is any small constant and the magnitude of t is either $\Omega(\sqrt{n})$ or $\Omega(n)$, depending on the properties of the games.[8] In particular, the authors prove the following results: If a game has low strategic sensitivity, then all its Nash equilibria are (t, ε)-fault-tolerant, with $t = \Omega(n)$; and if a game is anonymous, then it has at least one (t, ε)-fault-tolerant Nash equilibrium for which $t = \Omega(\sqrt{n})$.

A different interpretation of fault tolerance is the notion of viability, which is intended to capture the likelihood that an equilibrium arises or is maintained in any social or economic system. Kalai (2018) proposes two related indices to assess the viability of a Nash equilibrium σ. We focus on one of them here, the defection index $D(\sigma)$, defined as the smallest integer d for which the following holds: There are d players and a strategy for these players such that, when they play this strategy instead of σ, the strategy σ_i of some other player i is no longer a best response. Observe that a profile σ satisfies $D(\sigma) = d$ if and only if it is $(d - 1)$-fault-tolerant but not d-fault-tolerant.

[8]The asymptotic notation $\Omega(f(n))$ here means that there is some constant $c \in (0, 1]$ independent of n for which the actual value of t is at least $c \cdot f(n)$.

Based on observations from behavioral economics, Kalai (2018) argues that the defection index, along with other related indices, concisely captures the viability of equilibria: the greater the value of $D(\sigma)$, the more viable the equilibrium σ. If we consider an approximate version of this index—letting $D(\sigma)$ be the smallest d under which some other player's strategy is no longer an ε-best response—then the results of Gradwohl & Reingold (2014) imply that equilibria of large games are viable, with a defection index as high as $\Omega(\sqrt{n})$ or $\Omega(n)$, depending on the class of the game.

Since the symmetric SI game of Example 1 has low strategic sensitivity, the result of Gradwohl & Reingold (2014) implies that its equilibria are (t, ε)-fault-tolerant, with $t = \Omega(n)$. Hence, they are also viable, with defection index $\Omega(n)$. Additionally, since the asymmetric SI game of Example 3 is anonymous (but may not have low strategic sensitivity), it has at least one (t, ε)-fault-tolerant Nash equilibrium for which $t = \Omega(\sqrt{n})$. This equilibrium is thus also viable, with defection index $\Omega(\sqrt{n})$.

5. DIRECTIONS FOR FURTHER RESEARCH

5.1. The Independence Assumption

The assumption of independence imposed on the types of the players in the Bayesian game G limits the applicability of hindsight stability and related properties.

To see the difficulties associated with interdependent player types, consider the following modified incomplete-information SI game. On the day on which the game is to be played, a television program will be broadcast. The program's quality is unknown and is equally likely to be high (HQ) or low (LQ). There are two kinds of players: program-indifferent ones (INDIF), who are like the players in Example 2, with $h_i = 0.5$ and $o_i(p) = 1 - p$ (that is, they prefer to be out if and only if the proportion of people out is less than 0.5); and discriminating ones (DIS), who prefer to be out if and only if the program is of low quality. The distribution over each player's type is the same: The player is program indifferent with probability 1/3 and is a discriminator with probability 2/3. In the latter case, the player is told the realized quality of the program. The identical but not independent distribution over each player's type is thus uniform on INDIF, (DIS, HQ), and (DIS, LQ).

Notice that the discriminating types in this game have a dominant strategy: to be out if and only if they are told that the program's quality is LQ. This means that the only hindsight-stable outcome of any equilibrium occurs in the case in which no discriminating players are realized, a very unlikely event even for moderately large values of n.[9]

A natural and important direction for further research is to examine the extent and necessity of the independence limitation. One goal would be to identify properties of the interdependencies across types that capture real-world situations and under which robustness holds.

Some initial progress along these lines has been made. For example, Gradwohl & Reingold (2007) show that robustness holds in large games in which the types are only pairwise independent, while Gradwohl & Yehudayoff (2008) show the same for large games in which dependencies are captured by a dependency graph. Alternatively, one might turn to concentration-of-measure bounds that accommodate mild dependencies, such as those of Pelekis & Ramon (2017), and utilize them to derive hindsight stability.

[9]The only ε-hindsight-stable outcome of any equilibrium occurs when fewer than $2\varepsilon n$ discriminating players are realized. If more are realized, then with probability 1/2 the program is LQ and they are out, and with probability 1/2 the program is HQ and they are home. Indifferent players will have a deviation that increases their utility by at least ε in at least one of these cases.

A different approach is motivated by the observation that the lack of hindsight stability can be partially overcome when the game G, with its interdependent types, is played repeatedly. For instance, consider a variant of Example 2 in which players have interdependent types, and in which the SI game is played daily with these same types. Suppose that, at the end of every day, players obtain public information about the proportion of players who chose to be out that day. Finally, each player's utility for the day depends on the player's type, their chosen daily action, and the public information about that day's distribution of choices in the population.

For the case in which the number of players is large, Kalai & Shmaya (2018) develop a Markov-perfect imagined-continuum (MPIC) Nash equilibrium that captures the important strategic and informational aspects of the evolving equilibrium, despite the usual difficulties of updating private information in repeated games of this type. In a subsequent working paper, Kalai & Shmaya (2017) argue that with time the play of the MPIC equilibrium becomes predictable and hindsight stable, as if the players "learn to be independent." If we return now to our example with the television program, suppose the program is a television series with daily episodes, and such that the quality does not change much over time. After many daily repetitions of this SI game, the players will learn to predict the proportion of players who will choose to be out the next day, and at the end of that day no player would have a significant incentive to change their own choice even after observing others' actions.

A challenging important question for future research in this context is to address the issue of dynamically changing types. In the model of Kalai & Shmaya (2018), player types are fixed throughout the dynamic interaction. In many economic systems, however, player types change dynamically. For example, a player's type may describe resources that the player has available for that period's interaction in a market, or information that is available to the player but not to their opponents. Are there useful general models with changing types in which play converges to some kind of partial predictability and stability?

5.2. Game Embeddings

When studying a game, players and analysts often model the essentials of the game as if it were played in isolation. While such simplifications are useful and often unavoidable, in some situations they lead to distorted solutions. In the SI game, for example, if players are motivated to go out because they seek employment, the strategic interaction between employment seekers and employment providers should be included in the analysis.[10]

This example leads to a natural question: When is it proper to analyze a focal game G in isolation, despite its being embedded in a bigger metagame M? The structural robustness property described in the previous section provides a simple starting point.

Consider a pair (G, M) consisting of a focal game G and a metagame M, in which the latter satisfies the conditions ascribed in Section 4.3 to an alteration G^A of G. Consider any equilibrium profile σ of G and its adaptation σ^M to M. If all the G players play their σ_i^M strategies in M, then each one of them is maximizing their payoff no matter what strategies are chosen by non-G players of M. This is a desirable embedding property, since no strategies of the outsiders incentivize any of the G players to defect from their σ_i strategies.

This observation is useful from the perspective of both the players and the analysts attempting to predict the outcome of the interaction. From the players' perspective, if a game G is embedded

[10]The provision of financial assistance to unemployed workers in the United States may be viewed as a policy tool that disentangles such considerations.

in M, then the participants need not worry about the external effects of their actions in G. Instead, they have the option of using much simpler strategies, namely, the adaptation of their strategies from G to M. From an analyst's perspective, if one believes that players are more likely to choose simpler strategies, then it is possible to predict the outcome of the focal game G without worrying about the distortions introduced by M.

The interpretation above is limited by two issues: (*a*) the start of M, in which the types of the G players are assumed to be independent, and (*b*) the end of M, in which there is an assumption of payoff agreement between the G players' payoffs at the end of the metagame M with their payoffs in G. Both issues restrict the generality of metagames M in which G may be properly embedded. Independence rules out metagames in which the history of play that leads to the start of G creates dependencies among the types of the G players, and it is discussed in Section 5.1. Payoff agreement (specifically, the condition that playing G^A means playing G) rules out longer metagames in which the terminal payoffs at G are not the same as the continuation values in M. Removing this second limitation would be useful in broadening the applicability of game embeddings.

Repeated games may serve as a first step in tackling the payoff agreement issue. Consider a metagame M in which a game G is played twice: In the first stage, G^1, players play the game G; in the second stage, G^2, after being fully informed of the realized actions in G^1, the same players with the same realized types play G again. Suppose players' payoffs are the sums of their respective utilities in the two games. Here, one might wish to analyze the two games separately. However, extensive robustness does not follow immediately, even in large games, because payoff agreement can fail.

Suppose G^1 and G^2 are the incomplete-information SI games of Example 2, with the modification that players' types do not change in G^2. Consider the following embedded-games analysis: First, fix any equilibrium σ^1 of G^1. Then, suppose types are drawn again, and let σ^2 be any equilibrium of this modified variant of G^2. Then the hindsight stability of equilibria in G^1 and G^2 implies that the profile σ of the combined game, in which players play σ^1 in G^1 and then σ^2 in G^2, is an approximate equilibrium in G, even though types stay fixed after the first stage.

To see this, observe that, when n is sufficiently large, the fraction of players with types $h_i \geq C$ is close to C. Hence, for each σ^j, the fraction of realized players who are out must be close to C as well. It is straightforward to see that such a profile is an approximate equilibrium that is also hindsight stable.

One question for future research is: To what extent does this embedded-games analysis hold? Another question is the converse: Given any structurally robust, perfect equilibrium in the combined game, is it always the case that the strategies, when restricted to a stage game, also form an equilibrium?[11]

5.3. Aggregative Games

Finally, another direction for future research is a more thorough analysis of the general robustness properties of equilibria in aggregative games. As described in Section 3, such games are a generalization of anonymous games and have many applications in economics. Thus, questions about the strategic sensitivity of different classes of aggregative games (i.e., ones with different aggregation functions), and hence their robustness and stability properties, are important.

[11]This question is related to the work of Al-Najjar & Smorodinsky (2001) (building on Green 1982, Sabourian 1990), which shows that in large infinitely repeated games most players play myopically in each stage game.

SUMMARY POINTS

1. The Nash equilibria of games with many players are robust and stable.
2. They are hindsight stable, structurally robust, fault-tolerant, and viable.
3. Different classes of games—such as anonymous games, games with low strategic sensitivity, and general games—satisfy these robustness properties to varying degrees.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

The authors thank Adam Kalai for helpful discussions about the motivating examples. R.G. gratefully acknowledges the support of National Science Foundation award number 1718670.

LITERATURE CITED

Abraham I, Dolev D, Gonen R, Halpern J. 2006. Distributed computing meets game theory: robust mechanisms for rational secret sharing and multiparty computation. In *Proceedings of the Twenty-Fifth Annual ACM Symposium on Principles of Distributed Computing*, pp. 53–62. New York: ACM

Acemoglu D, Chernozhukov V, Werning I, Whinston MD. 2020. *Optimal targeted lockdowns in a multi-group SIR model*. NBER Work. Pap. 27102

Acemoglu D, Jensen MK. 2013. Aggregate comparative statics. *Games Econ. Behav.* 81:27–49

Adda J. 2016. Economic activity and the spread of viral diseases: evidence from high frequency data. *Q. J. Econ.* 131:891–941

Al-Najjar NI. 2008. Large games and the law of large numbers. *Games Econ. Behav.* 64:1–34

Al-Najjar NI, Smorodinsky R. 2001. Large nonanonymous repeated games. *Games Econ. Behav.* 37:26–39

Alon N, Spencer JH. 2004. *The Probabilistic Method*. Hoboken, NJ: Wiley

Alvarez FE, Argente D, Lippi F. 2020. *A simple planning problem for COVID-19 lockdown*. NBER Work. Pap. 26981

Aumann RJ, Shapley LS. 1974. *Values of Non-Atomic Games*. Princeton, NJ: Princeton Univ. Press

Azevedo EM, Budish E. 2019. Strategy-proofness in the large. *Rev. Econ. Stud.* 86:81–116

Azrieli Y, Shmaya E. 2013. Lipschitz games. *Math. Oper. Res.* 38:350–57

Babichenko Y, Barman S, Peretz R. 2014. Simple approximate equilibria in large games. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, pp. 753–70. New York: ACM

Blonski M. 1999. Anonymous games with binary actions. *Games Econ. Behav.* 28:171–80

Bodoh-Creed A. 2013. Efficiency and information aggregation in large uniform-price auctions. *J. Econ. Theory* 148:2436–66

Carmona G. 2008. Purification of Bayesian–Nash equilibria in large games with compact type and action spaces. *J. Math. Econ.* 44:1302–11

Carmona G, Podczeck K. 2012. Ex-post stability of Bayes–Nash equilibria of large games. *Games Econ. Behav.* 74:418–30

Carmona G, Podczeck K. 2020. Strict pure strategy Nash equilibria in large finite-player games. *Theor. Econ.* In press

Carroll G. 2019. Robustness in mechanism design and contracting. *Annu. Rev. Econ.* 11:139–66

Cartwright E, Wooders M. 2009. On equilibrium in pure strategies in games with many players. *Int. J. Game Theory* 38:137–53

Cerreia-Vioglio S, Maccheroni F, Schmeidler D. 2020. Equilibria of nonatomic anonymous games. arXiv:2005.01839 [econ.TH]

Che YK, Tercieux O. 2019. Efficiency and stability in large matching markets. *J. Political Econ.* 127:2301–42

Cole R, Tao Y. 2016. Large market games with near optimal efficiency. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pp. 791–808. New York: ACM

Daskalakis C, Papadimitriou CH. 2015. Approximate Nash equilibria in anonymous games. *J. Econ. Theory* 156:207–45

Deb J, Kalai E. 2015. Stability in large Bayesian games with heterogeneous players. *J. Econ. Theory* 157:1041–55

Eliaz K. 2002. Fault tolerant implementation. *Rev. Econ. Stud.* 69:589–610

Farboodi M, Jarosch G, Shimer R. 2020. *Internal and external effects of social distancing in a pandemic*. NBER Work. Pap. 27059

Gaffeo E. 2003. The economics of HIV/AIDS: a survey. *Dev. Policy Rev.* 21:27–49

Galeotti A, Goyal S, Jackson MO, Vega-Redondo F, Yariv L. 2010. Network games. *Rev. Econ. Stud.* 77:218–44

Geoffard PY, Philipson T. 1996. Rational epidemics and their public control. *Int. Econ. Rev.* 37:603–24

Goenka A, Liu L. 2012. Infectious diseases and endogenous fluctuations. *Econ. Theory* 50:125–49

Gradwohl R, Reingold O. 2007. *Partial exposure and correlated types in large games*. Work. Pap., Weizmann Inst. Sci., Rehovot, Isr.

Gradwohl R, Reingold O. 2010. Partial exposure in large games. *Games Econ. Behav.* 68:602–13

Gradwohl R, Reingold O. 2014. Fault tolerance in large games. *Games Econ. Behav.* 86:438–57

Gradwohl R, Yehudayoff A. 2008. t-Wise independence with local dependencies. *Inf. Process. Lett.* 106:208–12

Green EJ. 1982. Noncooperative price taking in large dynamic markets. In *Noncooperative Approaches to the Theory of Perfect Competition*, ed. A Mas-Colell, pp. 37–64. New York: Academic

Greenwood J, Kircher P, Santos C, Tertilt M. 2019. An equilibrium model of the African HIV/AIDS epidemic. *Econometrica* 87:1081–113

Jensen MK. 2010. Aggregative games and best-reply potentials. *Econ. Theory* 43:45–66

Jordan JS, Radner R. 1982. Rational expectations in microeconomic models: an overview. *J. Econ. Theory* 26:201–23

Kalai E. 2004. Large robust games. *Econometrica* 72:1631–65

Kalai E. 2008. Large games: structural robustness. In *The New Palgrave Dictionary of Economics*, ed. S Durlauf, LE Blume, pp. 920–32. Basingstoke, UK: Palgrave Macmillan. 2nd ed.

Kalai E. 2018. *Viable Nash equilibria: formation and defection.* Work. Pap., Northwest. Univ., Evanston, IL

Kalai E, Shmaya E. 2017. *Learning, predictions, and stability in large dynamic interactions.* Work. Pap., Northwest. Univ., Evanston, IL

Kalai E, Shmaya E. 2018. Large strategic dynamic interactions. *J. Econ. Theory* 178:59–81

Kearns M, Pai M, Roth A, Ullman J. 2014. Mechanism design in large games: incentives and privacy. In *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science*, pp. 403–10. New York: ACM

Keppo J, Quercioli E, Kudlyak M, Wilson A, Smith L. 2020. *The behavioral SIR model, with applications to the swine flu and COVID-19 pandemics*. Virtual Macro Pres., Apr. 10. **https://www.lonessmith.com/wp-content/uploads/2020/04/pandemic-slides.pdf**

Kermack WO, McKendrick AG. 1927. A contribution to the mathematical theory of epidemics. *Proc. R. Soc. A* 115:700–21

Khan MA, Rath KP, Yu H, Zhang Y. 2017. On the equivalence of large individualized and distributionalized games. *Theor. Econ.* 12:533–54

Khan MA, Sun Y. 2002. Non-cooperative games with many players. In *Handbook of Game Theory with Economic Applications*, Vol. 3, ed. R Aumann, S Hart, pp. 1761–808. Amsterdam: Elsevier

Lee S. 2016. Incentive compatibility of large centralized matching markets. *Rev. Econ. Stud.* 84:444–63

Lipton RJ, Markakis E, Mehta A. 2003. Playing large games using simple strategies. In *Proceedings of the 4th ACM Conference on Electronic Commerce*, pp. 36–41. New York: ACM

Milgrom PR, Weber RJ. 1985. Distributional strategies for games with incomplete information. *Math. Oper. Res.* 10:619–32

Pelekis C, Ramon J. 2017. Hoeffding's inequality for sums of dependent random variables. *Mediterr. J. Math.* 14:243

Rashid S. 1983. Equilibrium points of non-atomic games: asymptotic results. *Econ. Lett.* 12:7–10
Rubinstein A. 1991. Comments on the interpretation of game theory. *Econometrica* 59:909–24
Sabourian H. 1990. Anonymous repeated games with a large number of players and random outcomes. *J. Econ. Theory* 51:92–110
Schmeidler D. 1973. Equilibrium points of nonatomic games. *J. Stat. Phys.* 7:295–300
Wooders M, Cartwright E, Selten R. 2006. Behavioral conformity in games with many players. *Games Econ. Behav.* 57:347–60