Generalized Sampling Equilibrium

Josh Cherry and Yuval Salant

Northwestern University

December 3, 2014

*** WORK IN PROGRESS - PRELIMINARY AND INCOMPLETE ***

Abstract

We propose a solution concept called *Generalized Sampling Equilibrium* (GSE), where players use statistical rather than strategic reasoning. This concept is rooted in the sampling equilibrium of Osborne and Rubinstein (1998, 2003), and accommodates a variety of other statistical inference procedures. We show that the GSE is unique for a large class of two-action games, and characterize how it relates to the Nash equilibrium. We also characterize how the GSE changes with the size of the sample players obtain, and demonstrate the predictions of this solution concept in several applications including a labor matching environment. We show that sampling introduces a friction that results in larger unemployment than in Nash equilibrium.

1 Introduction

An individual's benefit from taking a particular action often depends on how many other individuals take that same action. For example, when deciding whether to attend some event, it is often important to know how many people plan to attend the event; When deciding whether to get vaccinated, the proportion of people who are or are not vaccinated against the disease is relevant; When a firm decides whether to enter a new market, the number of other potential entrants is important and so on.

In the standard game theoretic framework, rational players understand the structure of their environment and reason strategically about the behavior of others. According to this approach, every individual forms expectations about the proportion of people who will take each action; When everybody maximizes their utility given these expectations, the expectations are "correct" in the sense that the proportion of people who take each action is identical to the expectations.

Several alternatives to this approach have been discussed in the literature, the most relevant for the current paper being the Sampling Equilibrium of Osborne and Rubinstein (1998, 2003). In a sampling equilibrium, every individual "samples" either other individuals' actions or the (random) payoffs from his own actions. The individual then treats his sample as representative of the entire population to form point estimates about the proportion of people who take each action and maximizes his utility based on this (possibly, incorrect) belief. In equilibrium, the distribution from which every individual obtains his sample is a fixed point: if individuals sample from this distribution and maximize utility, then the proportion of individuals taking each action is identical to the distribution. Osborne and Rubinstein (2003) study the application of this solution concept to a voting model, in which each individual samples two or three other individuals.

Generalized Sampling Equilibrium, the solution concept introduced and studied in this paper, builds on Osborne and Rubinstein (2003) in two ways. First, rather than assuming agents' beliefs coincide with sample averages, we allow for individuals to understand that their sample may be noisy. Using only the sample average is one of many possible *reasoning procedures* players may use to form beliefs about the population with which to compute best replies. Our model accommodates, and compares the predictions of, a wide class of reasoning procedures. Second, we develop tools that operationalize sampling-based solution concepts to obtain existence, uniqueness and comparative statics in a variety of applications.

Related Literature

to be completed...

2 Model

A unit mass of players each decide whether or not to take an action A. Each player's utility from taking the action A is $u(\theta_i, \alpha) = \theta_i - f(\alpha)$, where θ_i is a player's type, i.e., his private benefit from taking the action A, and $f(\alpha)$ is the cost incurred by a player taking the action A if a proportion α of players take the action A. The benefit θ_i is distributed uniformly on $[0, 1]$, and the function f is a positive, increasing and continuous function with $0 \leq f(0), f(1) \leq 1$. That is, the highest type will take action A for any α, and there is a small enough α to make the lowest type weakly prefer action A. Assume that the utility from not taking the action A is normalized to 0.

In order to decide whether to take the action, each player has to reason about the proportion of players that will take action A. The standard game theoretic approach posits that every player forms the same "correct" belief about this proportion, in the sense that if every player maximizes utility given this belief, the proportion of players taking the action is identical to the belief. One way a player may arrive at the correct belief is by reasoning strategically about the situation. He may reason that when a proportion α plans to take the action A, all players with a type $\theta \leq f(\alpha)$ will take the action A. Thus the proportion has to be equal to $1 - f(\alpha)$. So the conjecture is correct when it satisfies $\alpha = 1 - f(\alpha)$. If players have common knowledge of a shared conjecture α^*_{NE}, then α^*_{NE} is Nash equilibrium. The proportion α^*_{NE} that solves $\alpha = 1 - f(\alpha)$ is the unique Nash equilibrium of the game.

Following Osborne and Rubinstein (1998, 2003), we propose an alternative solution to the game. According to this solution, each player in the population has a tentative plan whether to take the action or not. Players collect information by randomly asking k other players about their tentative plans. Players then infer from their random sample the proportion of players who tentatively plan on taking the action, and each reevaluates their own action based on their estimate. In a generalized sampling equilibrium, players sample from a *reliable source*: the proportion of players who initially plan to take the action is identical to the proportion of players who actually take the action after collecting information and potentially revising their planned action.

More formally, each player's tentative plan is either 1 or 0, depending on whether he plans to take the action or not. Let α denote the proportion in the population who tentatively plan to take the action. All players obtain a sample of k random draws from a Bernoulli random variable with a probability of success α. Given their sample, players form beliefs about the proportion α taking the action. A reasoning procedure describes their belief formation as a function of their sample size and their observed proportion of successes.

Definition (Reasoning Procedure)**.** *A **reasoning procedure** $G = \{G_z^k\}$ is a family of probability distributions over α indexed by the sample size $k \geq 1$ and the fraction of successes $z \in [0, 1]$ such that (1): Fixing the sample size k, $G_{z'}^k$ strictly first order stochastically dominates G_z^k for $z < z'$, and (2): Fixing the proportion of successes z, G_z^k is a mean-preserving spread of $G_z^{k'}$ for $k < k'$.*

A reasoning procedures satisfies two properties. The first is that fixing the sample size k, a player puts higher weight on more players taking the action for a higher observed proportion of successes. The second is that fixing the proportion of successes, the beliefs are less dispersed when the sample size is larger.

Note that a reasoning procedure is defined for any possible proportion of successes z, while the player actually only observes proportions on a $\frac{j}{k}$ grid. We need the additional structure when analyzing the structure of equilibria as the number of samples k, and therefore the grid, changes. Let $g_z^k(\alpha)$ be the density function of $G_z^k(\alpha)$ when it exists. Following are a few examples of reasoning procedures.

Example 2.1 (Sample Average, Osborne and Rubinstein (2003))**.** When a player observes a proportion z of successes, he reasons that a proportion z of players will take the action with probability 1, independently of his sample size. That is,

$$G_z^k(\alpha) = \begin{cases} 1 & \alpha \geq z \\ 0 & \alpha < z \end{cases}$$

$\diamond$

Example 2.2 (Bayesian Updating with a Uniform Prior)**.** Each player has a uniform prior over the proportion α, and uses Bayesian updating to obtain a posterior given a sample size k with a proportion z of successes. The density of the posterior belief given this reasoning procedure is given by

$$g_z^k(\alpha) = \frac{\alpha^{zk}(1-\alpha)^{(1-z)k}}{B(zk+1,(1-z)k+1)}$$

where $B(\cdot,\cdot)$ is the Beta function. Thus G_z^k is a Beta$(zk+1,(1-z)k+1)$ distribution. $\diamond$

Example 2.3 (Haldane's Reasoning Procedure)**.** Each player has complete ignorance about the proportion α taking the action. "Haldane's prior" (Haldane, 1932; Zhu and Lu, 2004), the limit of Beta(ϵ,ϵ) distributions as $\epsilon \to 0$, is often used in the statistics literature to capture such ignorance. After observing a sample size k with a proportion z of successes, the player updates his belief to the Beta$(zk,(1-z)k)$ distribution. Note that this reasoning procedure is non-Bayesian since there does not exist a proper prior that, together with Bayesian updating, generates a Beta$(zk,(1-z)k)$ posterior. $\diamond$

Comment on Beta Distributions. The Beta distribution is a conjugate prior for Binomial distributions. That is, given a Beta(a,b) prior, the posterior following the realization of a Binomial random variable with n draws is a Beta$(a+s,b+n-s)$ distribution, where s is the number of successes: The first argument is incremented by the number of successes, and the second incremented for the number of failures. The sum $a+b+n$ is the number of *pseudo-observations*, a way of measuring the relative weight placed on the prior $(a+b)$ compared to that on the sample n. A uniform prior on [0,1] corresponds to a Beta(1,1) distribution, or in other words, that after sampling the player holds beliefs that correspond to $k+2$ pseudo-observations, so he puts some weight on his prior. In contrast, a player using Haldane's reasoning procedure has no information, and bases his entire belief on the outcome of his sampling in the sense that the inferences drawn after observing a sample size k reflects k pseudo-observations.

Example 2.4 (Truncated Normal)**.** When a player observes a proportion z of successes in a sample of size k, he reasons that a proportion α will take the action, where α has a normal distribution with mean z and variance $\frac{z(1-z)}{k}$, truncated symmetrically. $\diamond$

Fixing a sample size k and a reasoning procedure, a type θ_i player who observed a proportion z successes in the sample best responds to his belief G_z^k. That is, he takes the action iff

$$\theta_i \geq F_k(z)$$

where $F_k(z)$ is the expectation of f under G_z^k

$$F_k(z) = \int_0^1 f(\alpha)\, dG_z^k(\alpha).$$

We now define the solution concept. Suppose a proportion α of players tentatively plans on taking action A, and each player observes the tentative plans of k other players. Then the probability of observing j successes in a sample of size k is

$$\binom{k}{j}\alpha^j(1-\alpha)^{k-j}.$$

Conditional on observing j successes, the proportion of players taking the actionis $\left(1-F_k\left(\frac{j}{k}\right)\right)$. Then the fraction of players observing j successes and taking action A is

$$\binom{k}{j}\alpha^j(1-\alpha)^{k-j}\left(1-F_k\left(\frac{j}{k}\right)\right).$$

Summing over j yields the total measure of players choosing action A

$$\sum_{j=0}^{k}\binom{k}{j}\alpha^j(1-\alpha)^{k-j}\left(1-F_k\left(\frac{j}{k}\right)\right).$$

In equilibrium, the aggregate tentative plan is correct in the sense that sampling from a population in which a proportion α plan to take the action A induces a proportion α to take the action A.

Definition. *A* Generalized Sampling Equilibrium *is a proportion* $\hat{\alpha}$ *such that*

$$\hat{\alpha}=\sum_{j=0}^{k}\binom{k}{j}\hat{\alpha}^j(1-\hat{\alpha})^{k-j}\left(1-F_k\left(\frac{j}{k}\right)\right).$$

A generalized sampling equilibrium makes weaker demands on players' knowledge and reasoning than Nash equilibrium. In terms of knowledge, all that a player needs to know (in addition to his available actions and utility) is that he obtains information from a reliable source in the sense that the tentative plan of the population reflects how the population will actually behave. In particular, knowledge of other players' incentives, rationality, information or how they reason about the games is not needed.

An Example: Consumption with Negative Externalities To illustrate the solution concept, consider a setting in which the consumption of a good by one player reduces the utility of consumption for all other agents. One example is visiting an amusement park: The more people that visit the park, the more congested it is, and hence the less enjoyable the experience is. Another example is purchasing a status good such as clothing item of a new style: As more individuals choose to own the good, the less effectively it conveys status.

Players derive a private benefit θ_i from consuming the good and suffer an externality cost of $\frac{1}{2}\alpha^2$. Specifically, we assume that the utility from consuming the good is $u(\theta_i, \alpha) = \theta_i - \frac{1}{2}\alpha^2$, i.e. the good is provided for free.[1] The Nash equilibrium proportion of individuals consuming the good is given by the solution to the equation

$$1 - \alpha = \frac{1}{2}\alpha^2$$

which has a unique solution on $[0, 1]$ at $\alpha = \sqrt{3} - 1 \approx .73$.

Suppose instead that each player samples the tentative plans of k other individuals and uses the sample average reasoning procedure to form beliefs about α. That is, a player who observes j out of k successes consumes the good if and only if the private benefit θ_i exceeds the cost $\frac{1}{2}\left(\frac{j}{k}\right)^2$.

Fix α. For a sample size $k = 1$, with probability α a player believes the entire population is consuming, and consumes only if his private benefit exceeds $\frac{1}{2}$. With probability $1 - \alpha$ a player believes no one is consuming the good, and since there is no externality, he consumes the good. Thus a GSE has to satisfy

$$\alpha = \alpha \cdot \frac{1}{2} + (1 - \alpha) \cdot 1.$$

This equation has a unique solution $\alpha = \frac{2}{3}$.

For $k = 2$, a GSE is characterized by

$$\alpha = \alpha^2 \cdot \frac{1}{2} + 2\alpha(1 - \alpha) \cdot \left(1 - \frac{1}{2}\left(\frac{1}{2}\right)^2\right) + (1 - \alpha)^2 \cdot 1$$

The equilibrium condition is quadratic in α, and has a unique positive root at $\alpha = \frac{1}{2}\left(\sqrt{41} - 5\right)$. In general, a generalized sampling equilibrium will be a proportion $\hat{\alpha}$ that solves the k-th order polynomial

$$\hat{\alpha} = \sum_{j=0}^{k} \binom{k}{j} \hat{\alpha}^j (1 - \hat{\alpha})^{k-j} \left(1 - \frac{1}{2}\left(\frac{j}{k}\right)^2\right)$$

Whether a solution in $[0, 1]$ exists or is unique is far from obvious.

3 Existence, Uniqueness, and Comparative Statics

In this section, we establish that the Generalized Sampling Equilibrium is unique for any reasoning procedure and for any number of samples. By adding the assumption on reasoning

[1] We extend this example to include prices in Section 4.1

procedures that mean beliefs equal sample averages, we can prove that the Nash equilibrium is higher (lower) than the GSE whenever f is convex (concave). We then introduce the second of two assumptions on reasoning procedures, which together allow us to provide comparative statics on how the unique GSE changes as the sample size gets larger. The proofs of these theorems, driven primarily by the connection between our equilibrium condition and Bernstein polynomials, can be found in the appendix.

Definition (Bernstein Polynomial)**.** *For any function $f(\alpha)$ defined on the closed interval $[0,1]$, the k-th* order Bernstein Polynomial *of $f(\alpha)$ is defined to be*

$$B_k(f;\alpha) \equiv \sum_{j=0}^{k} \binom{k}{j} \alpha^j (1-\alpha)^{k-j} f(j/k)$$

For any function f, the Bernstein Polynomial of order k approximates the function f at the $k+1$ points $\{0, \frac{1}{k}, \frac{2}{k}, \dots, \frac{k-1}{k}, 1\}$ with binomial weights prescribed by α. Thus $B^k(f;\alpha)$ is a function of $\alpha \in [0,1]$. In a GSE,

$$\begin{aligned}
\hat{\alpha} &= \sum_{j=0}^{k} \binom{k}{j} \hat{\alpha}^j (1-\hat{\alpha})^{k-j} \left(1 - F_k\left(\frac{j}{k}\right)\right) \\
1-\hat{\alpha} &= \sum_{j=0}^{k} \binom{k}{j} \hat{\alpha}^j (1-\hat{\alpha})^{k-j} \left(F_k\left(\frac{j}{k}\right)\right) \\
1-\hat{\alpha} &= B^k(F_k;\hat{\alpha})
\end{aligned}$$

Analyzing equilibria in our model in large part reduces to studying how the Bernstein operator acts on expected utility functions. We make most use of two types of results in the theory of Bernstein polynomials. The first is that the Bernstein operation maintains properties like the monotonicity and convexity of the function on which it operates. The second provides monotonicity properties of the operator's order k.

Theorem 1. *There exists a unique Generalized Sampling Equilibrium α_k^* for every number of samples $k \geq 1$.*

Proof of Theorem 1. In a Generalized Sampling Equilibrium

$$1-\alpha = B^k(F_k;\alpha)$$

$1-\alpha$ on the LHS is a strictly decreasing continuous function on $[0,1]$, which starts at 1 and ends at 0. The first order stochastic dominance of the reasoning procedure implies F_k is increasing, a

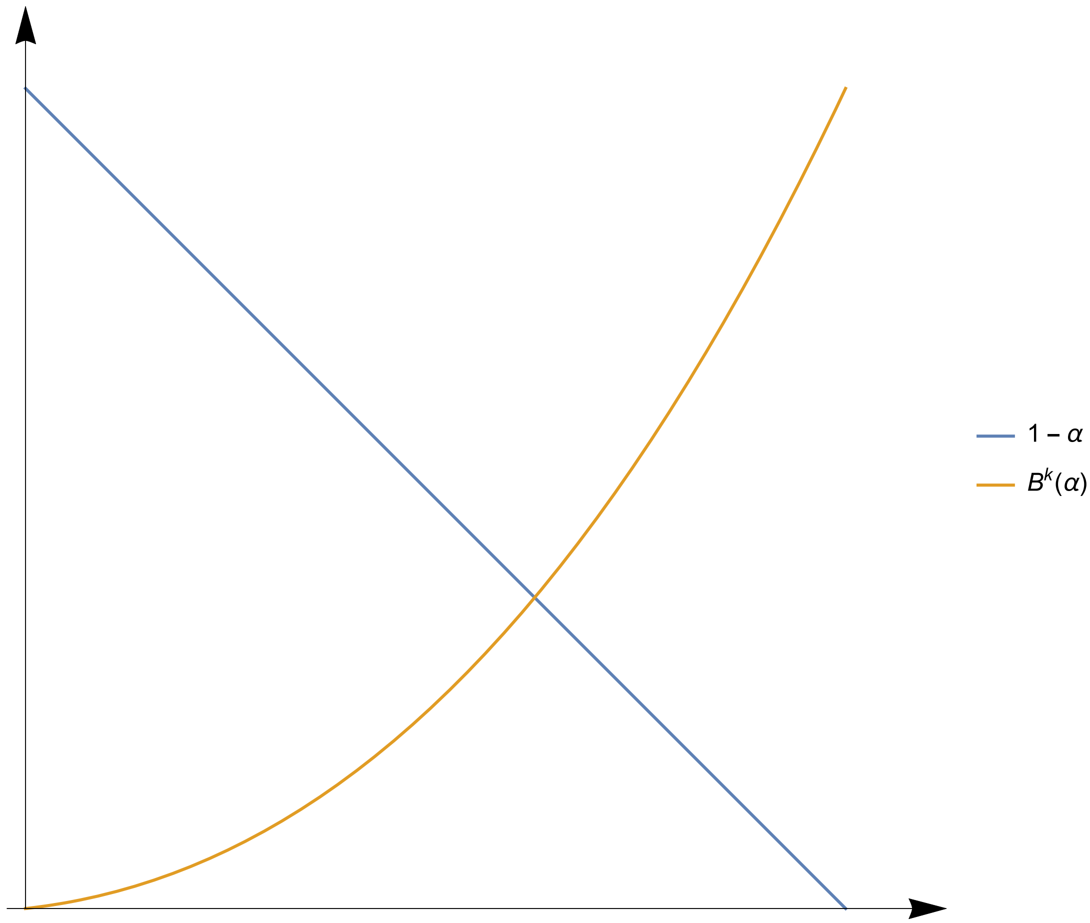

Figure 1: Generalized Sampling Equilibrium: Existence and Uniqueness

property inherited by the Bernstein polynomial $B^k(F_k;\alpha)$. The continuous function $B^k(F_k;\alpha)$ is increasing in α, and by definition equal to 0 at $\alpha = 0$ and equal to 1 when $\alpha = 1$. Thus the two functions cross exactly once. $\square$

Assumption 1 (Mean Preserving)**.** A reasoning procedure $G = \{G_z^k\}$ is *mean preserving* if, for every sample size k and any proportion of successes z, the expectation of α w.r.t. G_z^k is equal to z, i.e.

$$\int_0^1 \alpha \, dG_z^k(\alpha) = z$$

Assumption 1 is satisfied by the reasoning procedures in Examples 2.1, 2.3, and 2.4. However, the Bayesian reasoning procedure described in Example 2.2 fails to satisfy Assumption 1 because whenever a player has a proper prior and uses Bayesian updating, the posterior mean of α is a weighted average of the prior mean and the sample mean, where the weight is a function of k. Therefore it cannot coincide with the sample mean for all sample sizes.

Theorem 2. *If f is convex and the reasoning procedure is mean preserving, then the proportion α_k^* is strictly lower than α_{NE}^*. If f is concave, the ranking is reversed.*

Assumption 2 (Shape Preserving)**.** A reasoning procedure G is *shape preserving* if for every sample size k and f convex (concave), $F_k(z)$ is convex (concave) in z.

The reasoning procedure in Example 2.1 trivially satisfies Assumption 2 because $F_k = f$. A sufficient condition for Assumption 2 is that, fixing k, the reasoning procedure $G_z^k(\alpha)$ be mean preserving and totally positive of order 3 (TP3) in (z, α) (c.f. Jewitt (1988)).[2] The reasoning procedures in Examples 2.2, 2.3 and 2.4 all satisfy Assumption 2 as well since the distributions generated by those reasoning procedures belong to the exponential family, and exponential family densities are totally positive of all orders.

Given Assumptions 1 and 2, we can derive comparative statics of the equilibrium proportions taking action A as the number of samples k increases.

Theorem 3. *If f is convex, and the reasoning procedure is mean preserving and shape preserving, the proportion α_k^* is strictly lower than α_{k+1}^*, which in turn is strictly lower than α_{NE}^*. If f is concave, the reverse rankings hold.*

An immediate corollary of Theorem 3 relates any two mean and shape preserving reasoning procedures G, G' by the dispersion of the distributions they induce.

Corollary 1. *If f is convex and $G_k(z)$ is a mean-preserving spread of $G'_k(z)$ for all k, then $\alpha_k^* \leq \alpha_k'^*$. If f is strictly convex then $\alpha_k^* < \alpha_k'^*$. The inequalities are reverse if f is concave (resp. strictly concave).*

Proof. We prove the strictly convex case. First, observe that $F_k(z)$, the expectation of f under $G_k(z)$, is lower than $F'_k(z)$, the expectation of f under $G'_k(z)$, since f is strictly convex. Since both G and G' are shape preserving it follows, then, that $B^k(F_k(z); \alpha) < B^k(F'_k(z); \alpha)$ by the properties of Bernstein polynomials. $\square$

Another immediate corollary is that for any k, the reasoning procedure that assigns probability 1 to the sample average (as described in Example 2.1) obtains predictions that are closest

[2] A function $h(x, y)$ is totally positive of order 3 if for any $x_1 < x_2 < x_3$ and $y_1 < y_2 < y_3$ the matrix $(h(x_i, y_j)$ has a non-negative determinant for each minor of size ≤ 3. Total positivity of order 2 (TP2) is the Monotone Likelihood Ratio Property, and is implied by TP3. Therefore TP3 ensures that the likelihood ratios increase sufficiently quickly to preserve convexity under integration.

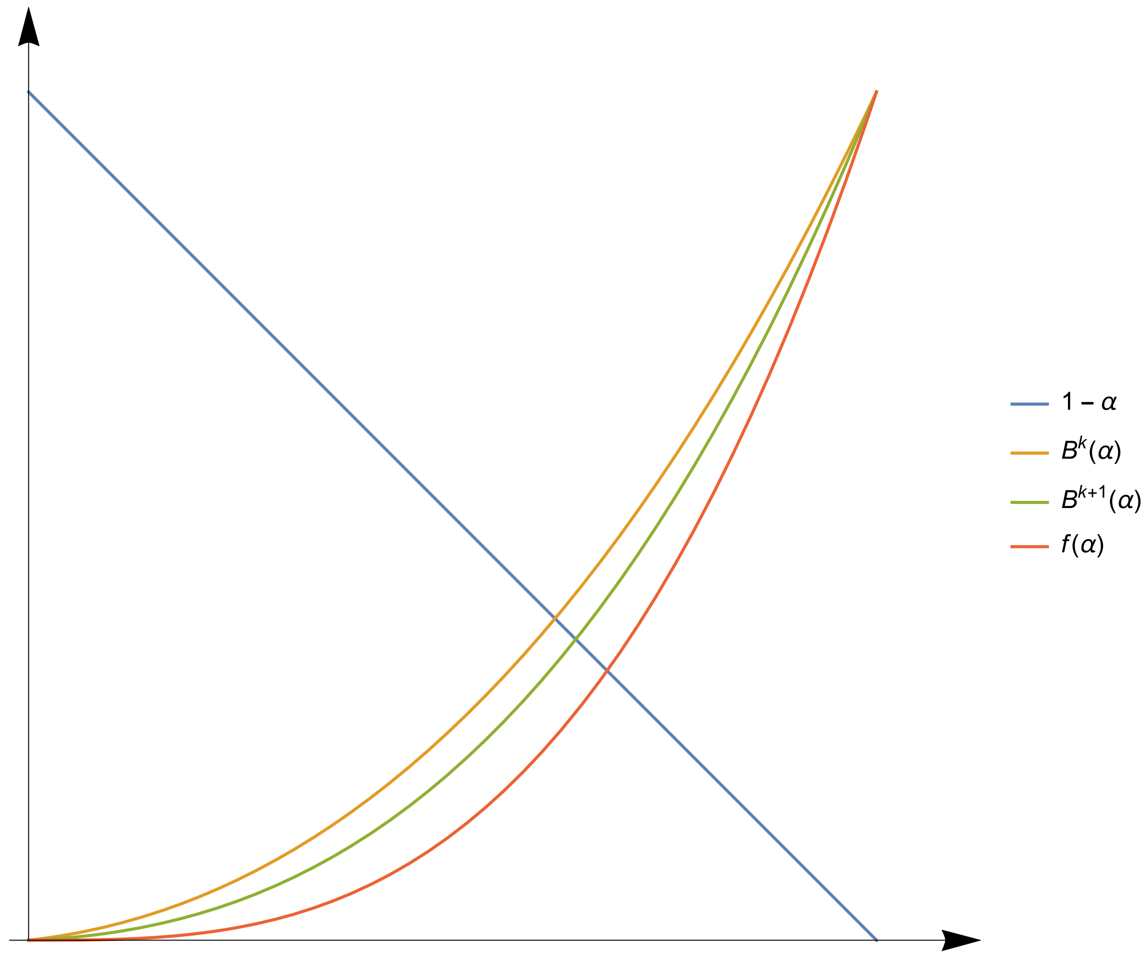

Figure 2: Generalized Sampling Equilibrium: Existence and Uniqueness

to the Nash equilibrium. We therefore focus on this reasoning procedure in what follows, as it provides a lower bound on the difference between GSE and Nash equilibrium.

Comment on Convergence. Osborne and Rubinstein (1998) proves that as $k \to \infty$ the sequence of sampling equilibria converges to a Nash equilibrium. Allowing k to grow without bound is not sufficient to guarantee convergence to Nash equilibrium in a Generalized Sampling Equilibrium. To see this, consider a mean preserving and shape preserving reasoning procedure G that has noise in the limit, i.e. there exists $\epsilon > 0$ such that the variance of $G_z^k \geq \epsilon$ for all k. Thus when G_z^k does not collapse to the point mass H_z as k tends to infinity, GSE will be bounded away from the Nash equilibrium. This observation suggests a sufficient condition for convergence of GSE to Nash equilibrium. Call a reasoning procedure $G = \{G_z^k\}$ *noiseless in the limit* if $G_z^k \xrightarrow{d} H_z$ for all z.

Corollary 2. *If a reasoning procedure G is noiseless in the limit, then* $\lim_{k\to\infty} \alpha_k^* = \alpha_{NE}^*$.

Intuitively, it is not enough that players collect enough information to make suitable inferences, their reasoning procedure must actually make use of a law of large numbers. $\diamond$

We will now illustrate the implications of Theorems 1, 2 and 3 in two specific settings.

Labor Supply

A unit mass of small firms use capital and labor to produce a consumption good x using the Cobb-Douglas production function $x(K, L) = (L + c)^{\gamma} K^{1-\gamma}$ for $c > 0$.[3] Suppose each firm has a single, fixed unit of capital. Then the only variable input is labor, which contributes to output with decreasing returns $x(L) = (L+c)^{\gamma}$. Assume for simplicity that the demand for x is perfectly elastic at a price normalized to 1.

Profit maximization by a firm implies that it will continue hiring labor until the the marginal product of labor $x'(L) = \gamma(L + c)^{\gamma-1}$ equals the market wage. Since the firms are symmetric they all employ the same amount of labor. If L units of labor are supplied inelastically, the labor market will clear at a wage $w(L) = \gamma(L + c)^{\gamma-1}$.

There is a unit mass of players, and each of them needs to decide whether or not to supply a unit of labor to the market. Each individual's private benefit of leisure is θ_i, and he has to give up this private benefit if he decides to work. Each individual's wage $w(L) = \gamma(L + c)^{\gamma-1}$ depends on the proportion of individuals supplying labor. Thus, an individual will supply labor only if he thinks $w(L)$ is larger than θ_i.

In order to estimate $w(L)$, each individual samples other workers' labor supply plans. An individual obtaining a sample of j successes infers that a fraction $\frac{j}{k}$ of potential workers will choose to supply a unit of labor, and that the wage will be $w\left(\frac{j}{k}\right)$. Therefore the individual will choose to work if $w\left(\frac{j}{k}\right)$ exceeds his private benefit of leisure θ_i.

To establish how the proportion of workers in a GSE relates to that in a Nash equilibrium, and how the proportion changes with the sample size, we need to rewrite utilities in terms of "not supplying labor". Let $N = 1 - L$ denote the proportion of individuals not supplying labor and let $z(N) = w(1-L)$ be the wage if a proportion N do not work. The utility of not supplying labor (relative to supplying labor) is then $\theta_i - z(N)$. Since $w(L) = \gamma(L+c)^{\gamma-1}$ is decreasing and convex, $z(N)$ is increasing and concave, so by Theorems 2 and 3 the proportion of individuals

[3] We assume that firms can still generate positive output in the absence of labor to ensure that the marginal product of labor (and therefore the wage) at $L = 0$ is defined.

not supplying labor is decreasing in the sample size and above the Nash equilibrium proportion. Therefore, the proportion of individuals supplying labor increases in the sample size and is below the Nash equilibrium proportion.

Demand for Vaccinations

A unit mass of players face the potential to contract and spread an infectious disease such as the flu. A vaccine exists that fully protects recipients from infection. Individuals who are not vaccinated may contract the disease one of two ways. When flu season begins, a proportion c of unvaccinated individuals contract the disease. In subsequent time periods, an agent who has not previously been infected may contract the flu by contact with an infected individual. An infected individual is contagious for 1 period. The utility of a player who remains healthy is 0, while getting sick delivers a payoff of -1. Each individual has a private cost θ_i of getting vaccinated, reflecting heterogeneous costs of traveling to a clinic or varying sensitivities to the vaccination's side effects.

For simplicity, suppose there are just three time periods. In period 0, individuals simultaneously decide once-and-for-all whether to obtain a vaccination. Those not vaccinated face a chance c of being infected. Thus if proportion α choose vaccination, then the mass of infected individuals at the end of period 0 is $c\,(1-\alpha)$.

In subsequent periods, individuals meet each other randomly. When an agent previously uninfected meets an infected agent, he becomes infected. Denote by i_t the mass of *infected* individuals at the end of period t, s_t the mass of *susceptible* (i.e. previously uninfected) individuals at the end of period t, and r_t is the mass of *removed* (i.e. vaccinated or recovered) individuals at the end of period t.

The environment evolves according to the following equations:

$$i_t = s_{t-1} i_{t-1}$$

$$s_t = s_{t-1}(1 - i_{t-1})$$

$$r_t = r_{t-1} + i_{t-1}$$

with the initial conditions:

$$i_0 = c\,(1-\alpha)$$

$$s_0 = 1 - \alpha - c\,(1-\alpha)$$

$$r_0 = \alpha$$

At the end of period 2, the total mass of agents infected at some point $I(\alpha; c)$ is

$$I(\alpha; c) = i_0 + i_1 + i_2$$

Substituting the expressions that describe the evolution of the system we obtain

$$I(\alpha; c) =$$

$$(1-\alpha)\, c\, \big((\alpha-3)\alpha - (1-\alpha)^3 c^3 - (2\alpha-3)(1-\alpha)^2 c^2 - (\alpha-2)^2(1-\alpha)c + 3\big)$$

This can be simplified to[4]

$$I(\alpha; c) = \gamma_0 + \gamma_1\alpha + \gamma_2\alpha^2 + \gamma_3\alpha^3 + \gamma_4\alpha^4$$

The ex-ante expected infection probability for an individual remaining unvaccinated is decreasing and convex in the proportion of others obtaining vaccinations.[5] Recasting the problem so the action is "not vaccinated" and denoting the proportion not obtaining the vaccination by $\rho = 1 - \alpha$, a type θ_i player who observed j out of k people tentatively planning on not being vaccinated, will not obtain the vaccination if and only if $\theta_i \geq I(1 - \frac{j}{k}; c)$. A proportion ρ abstaining from vaccination constitutes a GSE if

$$1 - \rho = \sum_{j=0}^{k} \binom{k}{j} \rho^j (1-\rho)^{k-j} I\left(1 - \frac{j}{k}; c\right)$$

will get vaccinated. By Theorems 2 and 3, the GSE rate of abstention is lower than in the Nash equilibrium and is increasing in the sample size. Therefore the vaccination rate is higher than that in the Nash equilibrium and is decreasing in the sample size k.

When players use Haldane's reasoning procedure (as described in Example 2.3), rather than the sample average, Corollary 1 indicates that the vaccination rate is strictly between the Nash equilibrium vaccination rate and the GSE vaccination rate with the sample average reasoning procedure.

[4]where the constants γ_i are $\gamma_0 = -c^4 + 3c^3 - 4c^2 + 3c$, $\gamma_1 = 4c^4 - 11c^3 + 12c^2 - 6c$, $\gamma_2 = -6c^4 + 15c^3 - 13c^2 + 4c$, $\gamma_3 = 4c^4 - 9c^3 + 6c^2 - c$, and $\gamma_4 = -c^4 + 2c^3 - c^2$.

[5]It is easily verified that $I(\alpha; c)$ is decreasing and convex in α, and increasing and concave in c for all $0 \leq \alpha \leq 1$ and $0 \leq c \leq 1$. Similarly, one can show that the probability of eventual infection is increasing and concave in the initial infection rate c for a given proportion α being vaccinated.

4 Applications

4.1 Selling to Consumers who Sample

Returning to the consumption with negative externalities example above, suppose now that a profit-maximizing monopolist sets a price p that consumers observe before deciding whether to obtain the good. For simplicity, assume the monopolist faces zero marginal cost. A consumer's utility from purchasing the good is $u(\theta_i, p, \alpha) = \theta_i - p - \frac{1}{2}\alpha^2$. After observing the price p but prior to acting, players sample the tentative plans of other potential customers. A player observing j out of k successes attends if and only if

$$\theta_i \geq p - \frac{1}{2}\alpha^2$$

Fix the price $p \leq \frac{1}{2}$ for the monopolist.[6] Then, a GSE in the subgame requires

$$\alpha = \sum_{j=0}^{k} \binom{k}{j} \alpha^j (1-\alpha)^{k-j} \left(1 - p - \frac{1}{2}\left(\frac{j}{k}\right)^2\right)$$

and by Theorem 1, there is a unique $\alpha_k(p)$ that solves this equation.

Fixing k, $\alpha_k(p)$ is decreasing in p. We thus have for every k a well-defined demand function. Rewrite the expression for a GSE as the inverse demand

$$p(\alpha) = 1 - \alpha - B_k(f; \alpha)$$

Since f is convex, the properties of the Bernstein Polynomials imply that $p(\alpha)$ is decreasing and concave.

Next, we turn to the monopolist's problem. The monopolist seeks to solve

$$\max_{\alpha} \Pi(\alpha) = \alpha \cdot p(\alpha) = \alpha\,(1 - \alpha - B_k(f; \alpha))$$

Since $p(\alpha)$ is concave, so is the objective function, and the first order condition

$$MR_k(\alpha_k^*) \equiv 1 - 2\alpha_k^* - B_k(f; \alpha_k^*) - \alpha_k^* B_k'(f; \alpha_k^*) = 0$$

is sufficient for optimality. For a sample of size k, let $p_k(\alpha_k^*)$ denote the profit maximizing price, let α_k^* be the profit maximizing quantity (i.e. proportion of customers served at the optimum), and let $\Pi_k^* = \Pi(\alpha_k^*)$ denote the monopolist's optimal profit.

Theorem 4. *The monopolist's profit Π_k^*, quantity α_k^*, and price $p_k(\alpha_k^*)$ are all increasing in the sample size k.*

[6] We prove in what follows that the monopolist would never choose a price above $\frac{1}{2}$.

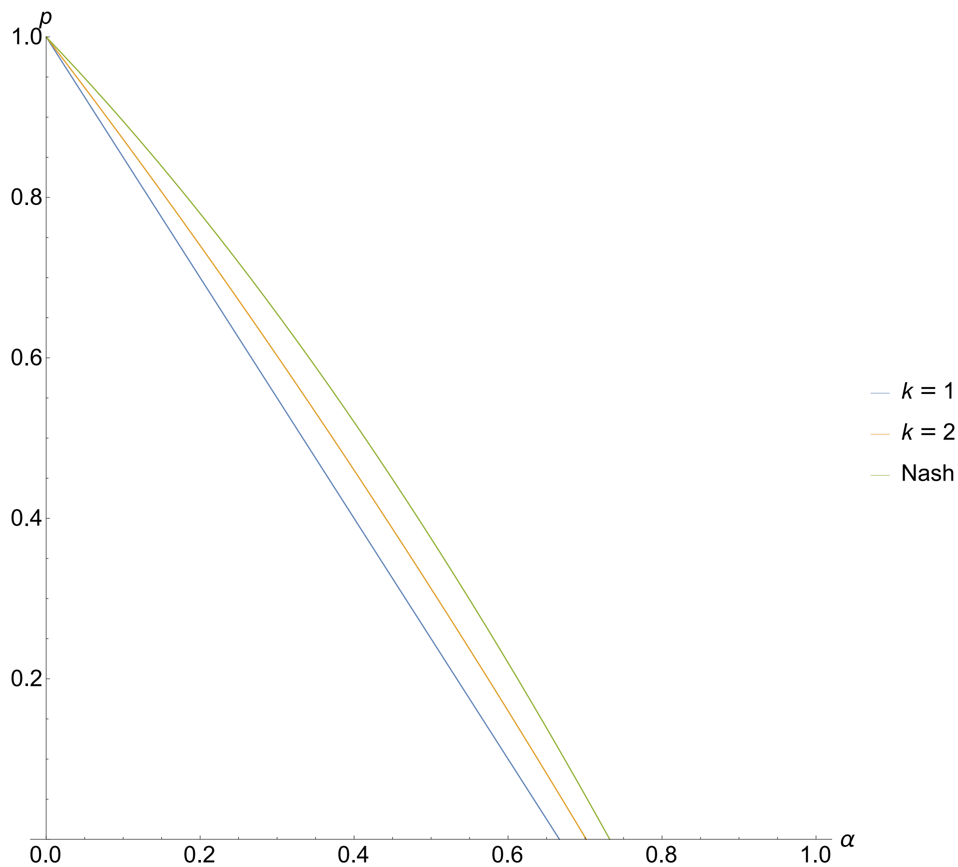

Figure 3: Demand Curves Faced by the Monopolist

4.2 Labor Matching Market

In this market, there are many workers and firms who must engage in costly search to create employment outcomes. Assume that a unit mass of workers seek job vacancies at a unit mass of firms. When a proportion α of workers choose to search for jobs and a proportion β of firms advertise positions, the number of vacancies filled is given by:

$$m(\alpha, \beta) \equiv \mu \alpha^x \beta^y$$

Assume $x + y = 1$, so that the matching function $m(\cdot, \cdot)$ exhibits constant returns to scale.[7] Further, we assume $0 < \mu < 1$, i.e. the presence of labor market frictions that would prevent full employment even if all workers and all firms participated. Denote by $g(\alpha, \beta) \equiv \frac{m(\alpha,\beta)}{\alpha}$ the utility of a worker when a proportion α of workers participate and a proportion β of firms participate. Analogously, denote by $h(\alpha, \beta) \equiv \frac{m(\alpha,\beta)}{\beta}$ the utility of a firm when a proportion α of workers participate and a proportion β of firms participate. Normalize to 0 the payoff to abstaining from the market.

Prior to the market, each worker must decide whether or not to participate in the market. Participation requires a worker to incur a cost θ_i, perhaps for preparing a resume', purchasing interview clothing, etc. Similarly, each firm must decide whether or not to participate in the market. Participation also requires a firm to incur a cost ω_i, perhaps for advertising the vacancy, administrative efforts to allocate the salary in the budget, etc.

Before making participation decisions, workers are able to sample the tentative participation plans of k firms. Upon observing j out of the k firms intending to participate, a worker believes a fraction $\beta = \frac{j}{k}$ will participate in aggregate. A worker with cost θ_i believes the utility of going to the labor market is positive iff

$$\mu \frac{\left(\frac{j}{k}\right)^y}{\alpha^{1-x}} \geq \theta_i$$

Define the cost of a worker who is indifferent between participation and abstention by

$$\bar{\theta}_j \equiv \mu \frac{\left(\frac{j}{k}\right)^y}{\alpha^{1-x}}$$

In the same manner, we can describe the behavior of firms, who are able to observe a sample of

[7] As noted by Pissarides (2000, page 4) the matching function summarizes a trading technology between heterogeneous agents that is not made explicit."

k workers' tentative plans. A firm with cost ω_i believes the utility of participation is positive iff

$$\mu \frac{\left(\frac{j}{k}\right)^x}{\beta^{1-y}} \geq \omega_i$$

Define the cost of a firm that is indifferent between participation and abstention by

$$\bar{\omega}_j \equiv \mu \frac{\left(\frac{j}{k}\right)^x}{\beta^{1-y}}$$

Notice that since $0 < x, y < 1$, both $\bar{\theta}_j$ and $\bar{\omega}_j$ are concave in j. Now suppose the tentative plans of workers and firms are α and β, respectively. Then the probability that a worker (firm) observes j out of k successes is $\binom{k}{j}\beta^j(1-\beta)^{k-j}$ (resp. firms $\binom{k}{j}\alpha^j(1-\alpha)^{k-j}$), and of those, a proportion $\bar{\theta}_j$ (resp. firms $\bar{\omega}_j$) choose to participate. To obtain the total number of workers and firms participating, we sum over j on both sides of the market and substitute the definitions of $\bar{\theta}_j$ and $\bar{\omega}_j$. This yields the system

$$\alpha = \frac{\mu}{\alpha^{1-x}} \sum_{j=0}^{k} \binom{k}{j} \beta^j (1-\beta)^{k-j} \left(\frac{j}{k}\right)^y \tag{1}$$

$$\beta = \frac{\mu}{\beta^{1-y}} \sum_{j=0}^{k} \binom{k}{j} \alpha^j (1-\alpha)^{k-j} \left(\frac{j}{k}\right)^x \tag{2}$$

which can be rearranged into

$$\alpha(\beta) \equiv \left(\mu \sum_{j=0}^{k} \binom{k}{j} \beta^j (1-\beta)^{k-j} \left(\frac{j}{k}\right)^y\right)^{\frac{1}{2-x}} \tag{3}$$

$$\beta(\alpha) \equiv \left(\mu \sum_{j=0}^{k} \binom{k}{j} \alpha^j (1-\alpha)^{k-j} \left(\frac{j}{k}\right)^x\right)^{\frac{1}{2-y}} \tag{4}$$

By the properties of Bernstein polynomials, the system describes (for every number of samples k) a function $\alpha_k(\beta)$ that is concave in β and a function $\beta_k(\alpha)$ that is concave in α.

Definition (Labor Matching Equilibrium)**.** *An equilibrium in the labor matching market is a pair of proportions (α_k^*, β_k^*) of workers and firms satisfying the system*

$$\alpha_k(\beta_k^*) = \alpha_k^*$$

$$\beta_k(\alpha_k^*) = \beta_k^*$$

Theorem 5. *For every number of samples k, there exists a unique Labor Matching Equilibrium with positive employment. Labor Market participation and equilibrium employment are increasing in k, and strictly below the Nash equilibrium level.*

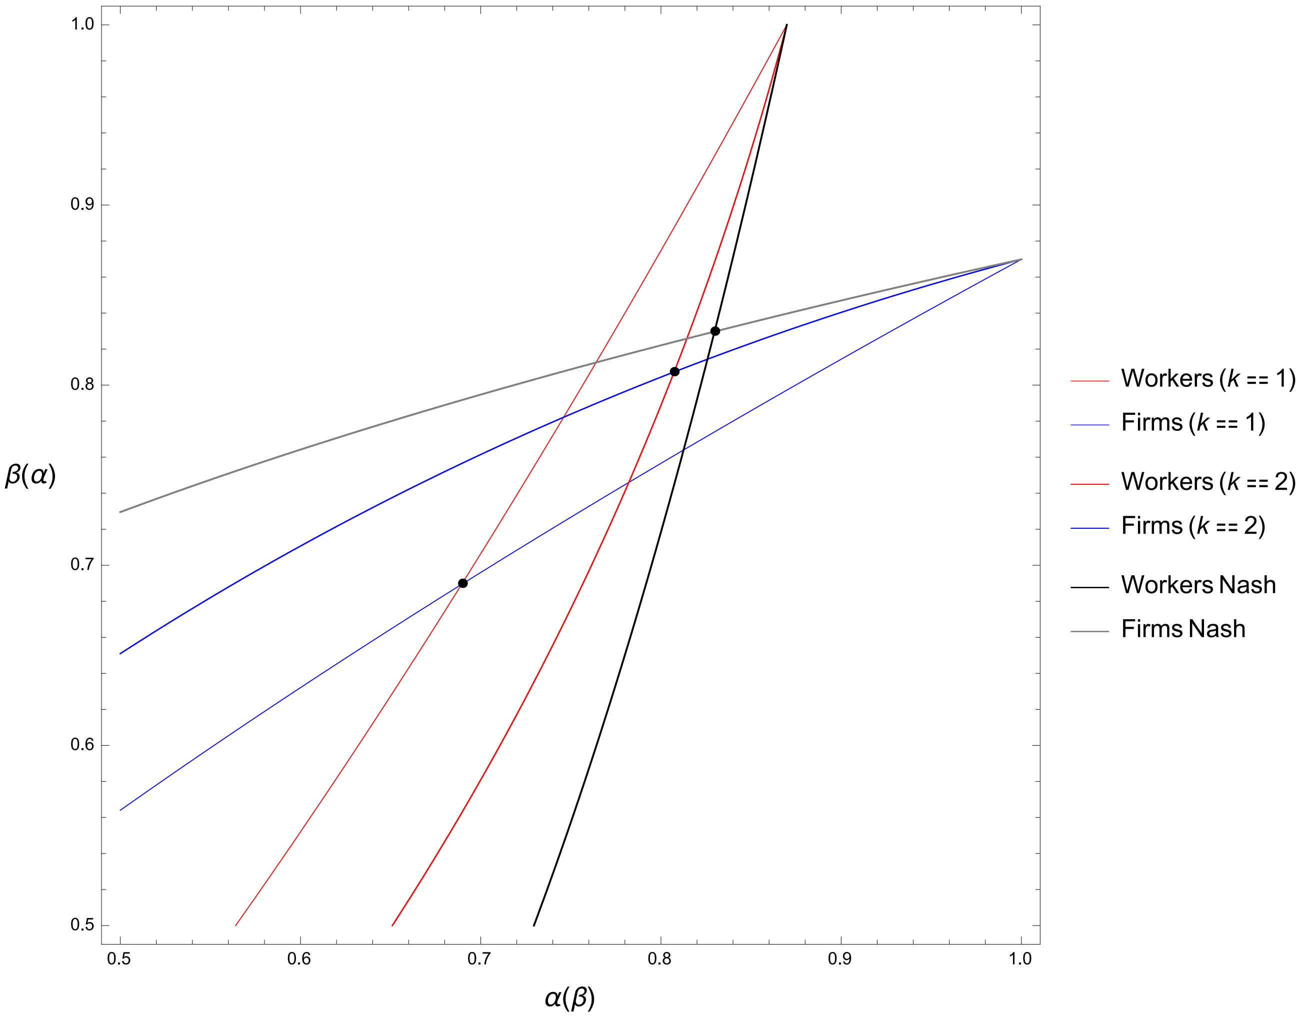

Figure 4: Labor Matching Equilibrium: $\mu = \frac{9}{10}$, $x = y = \frac{1}{2}$

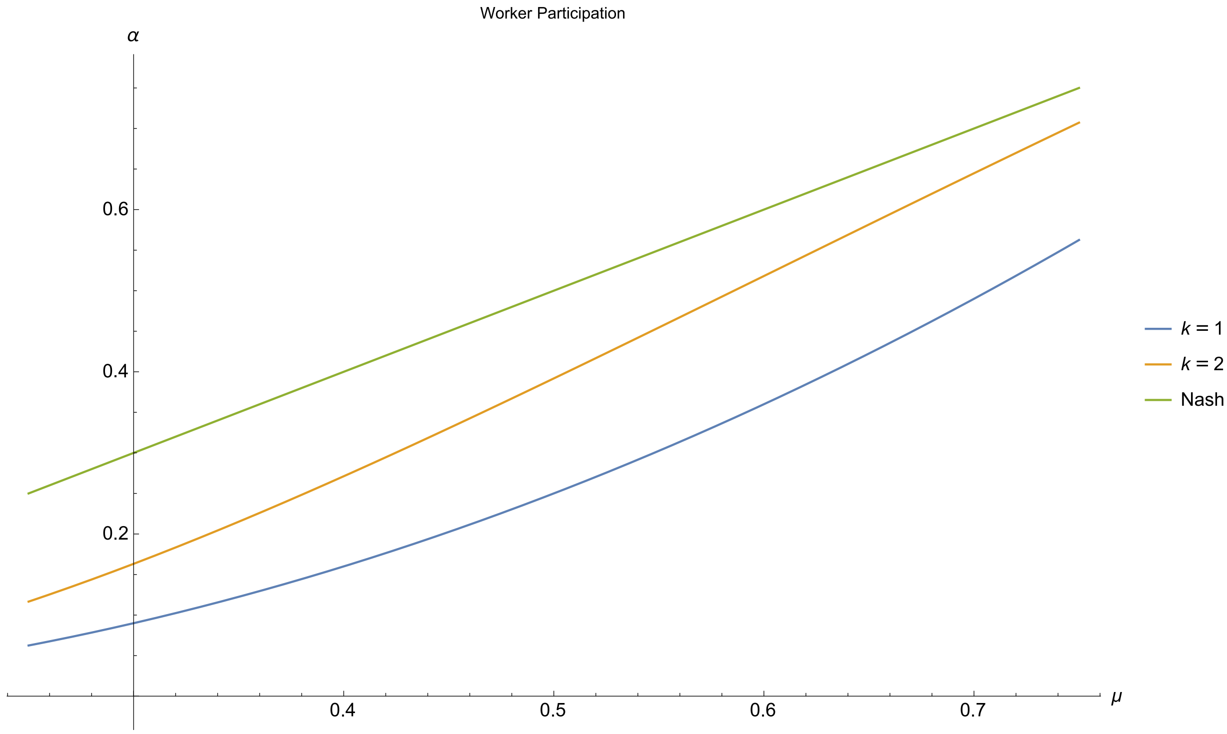

Figure 5: α as a function of μ

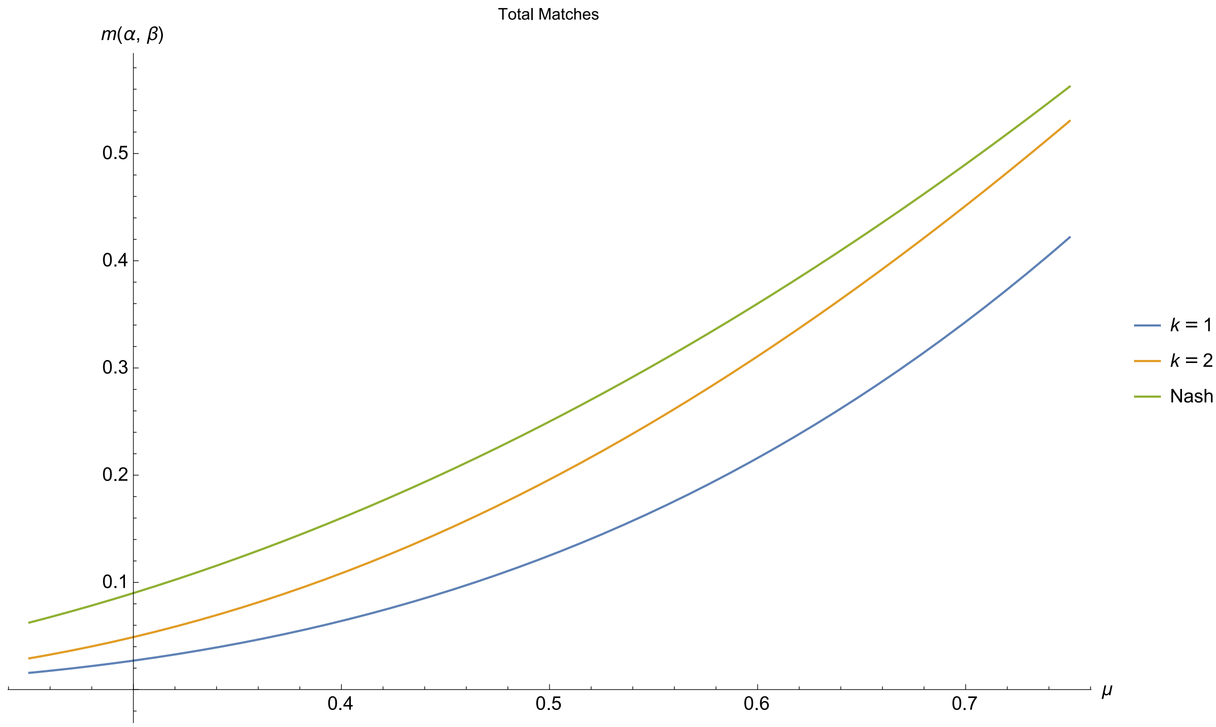

Figure 6: Total Matches as a function of μ

Appendix

Proof of Theorems 2 & 3. We prove the theorems for a convex function f. The proof for a concave function is analogous. It is sufficient (c.f. Figure 2) to prove that for all $\alpha \in [0, 1]$

$$B^k(F_k; \alpha) \geq B^{k+1}(F_{k+1}, \alpha) \geq f(\alpha)$$

We establish these inequalities in three steps. Note that only the first two steps are necessary for proving Theorem 2.

Step 1. $F_k(z) \geq F_{k+1}(z) \geq f(z)$:

Since f is convex, the reasoning procedure property that $G_z^k \prec^{SOSD} G_z^{k+1}$ implies, by Jensen's Inequality, that $F_k(z) \geq F_{k+1}(z)$. Let H_z be the distribution that places probability 1 on z: $H_z(\alpha) = 0$ if $\alpha < z$ and $H_z(\alpha) = 1$ if $\alpha \geq z$. Then $F_k(z) \geq f(z)$ by Jensen's Inequality, since $G_z^k \prec^{SOSD} H_z$. Appealing to this argument pointwise in z as it ranges from 0 to 1 completes this step. $\triangleleft$

Step 2. $B^k(F_k; \alpha) \geq B^k(F_{k+1}; \alpha) \geq f(\alpha)$

We invoke two monotonicity properties of Bernstein Polynomials. First, if v, w are real functions on $[0,1]$ and $v \geq w$, then for all $k \geq 1$

$$B^k(v;\cdot) \geq B^k(w;\cdot)$$

Second, if $v(\alpha)$ is convex, then for all $k \geq 1$

$$B^k(v;\alpha) \geq B^{k+1}(v;\alpha) \geq v(\alpha), \quad \alpha \in [0,1]$$

These properties, together with the inequality established in Step 1, prove this step. $\triangleleft$

Having established $B^k(F_k;\alpha) \geq f(\alpha)$, Theorem 2 is proved.

Step 3. $B^k(F_k;\alpha) \geq B^{k+1}(F_{k+1};\alpha)$ The reasoning procedure is shape preserving, so F_k and F_{k+1} are convex. We again appeal to the fact that Bernstein polynomials overestimate convex functions to obtain

$$B^k(v;\alpha) \geq B^{k+1}(v;\alpha) \geq v(\alpha)$$

Taking $v = F_{k+1}$ implies that $B^k(F_{k+1};\alpha) \geq B^{k+1}(F_{k+1};\alpha)$ $\triangleleft$

Combining the above results yields $B^k(F_k;\alpha) \geq B^{k+1}(F_{k+1},\alpha) \geq f(\alpha)$. $\square$

Proof of Theorem 4. To establish later results, we first show that the firm will never choose a quantity strictly greater than 1/2.

Proof. The result is proven if the marginal revenue of the firm is negative at $\alpha = 1/2$. The marginal revenue is

$$MR_k(1/2) = 1 - 2(1/2) - B_k(f;1/2) - 1/2B_k'(f;\alpha_k^*) < 0$$

since $B_k(f;\alpha)$ and $B_k'(f;\alpha)$ are both positive for all α, marginal revenue at 1/2 is negative. $\square$

Observation. *The monopolist's profit is increasing in k, and is largest in the Nash equilibrium.*

Proof. Appealing to Theorem 3, as we range over prices $p \in [0,\frac{1}{2}]$, we obtain $\alpha_k(p) \leq \alpha_{k+1}(p) \leq \alpha(p)$ where $\alpha(p)$ is demand in the rational case. The monopolist faces higher demand at every price, and thus earns weakly higher profits. $\square$

We are now in a position to state our first comparative static, namely that as the number of samples consumers take grows, the equilibrium quantity sold by the firm increases monotonically.

Lemma 1. *For all* $k \geq 1$

$$\alpha_k^* < \alpha_{k+1}^* < \alpha_R^*$$

Proof. To prove $\alpha^*_k \leq \alpha^*_{k+1}$ it suffices to show that if $MR_k(\alpha^*_k) = 0$ then $MR_{k+1}(\alpha^*_k) > 0$, for then the monopolist has a strict incentive to raise output. Thus we compare the marginal revenue with k samples to that with $k+1$ samples, noting that by the previous lemma $\alpha^*_k \leq 1/2$:

$$MR_k(\alpha^*_k) = 1 - 2\alpha^*_k - B_k(F_k; \alpha^*_k) - \alpha^*_k B'_k(F_k; \alpha^*_k) = 0$$

$$MR_{k+1}(\alpha^*_k) = 1 - 2\alpha^*_k - B_{k+1}(F_{k+1}; \alpha^*_k) - \alpha^*_k B'_{k+1}(F_{k+1}; \alpha^*_k)$$

Since $B_k(F_k; \alpha^*_k) < B_{k+1}(F_{k+1}; \alpha^*_k)$ for $\alpha^*_k \in [0, 1]$ and $B'_k(F_k; \alpha^*_k) < B'_{k+1}(F_{k+1}; \alpha^*_k)$ for $\alpha^*_k \in [0, 1/2]$, it follows that $MR_{k+1}(\alpha^*_k) > 0$.

A similar argument establishes that $\alpha^*_k \leq \alpha^*_R$ for all $k \geq 1$ by noting that by a property of the Bernstein Polynomials $B_k(f; \alpha) > f(\alpha)$ for all $\alpha \in [0, 1]$ and since $f(\alpha)$ is convex, $B'_k(F_k; \alpha) > f'(\alpha)$ for all $\alpha \in [0, 1/2]$. It follows immediately that the marginal revenue for a monopolist facing rational consumers at the optimal quantity for k-sample boundedly rational consumers is positive. □

□

References

Adell, J., Germán Badía, F., and de la Cal, J. (1993). Beta-type operators preserve shape properties. *Stochastic Processes and their Applications*, 48(1):1–8.

Al-Najjar, N. I. (1995). Decomposition and Characterization of Risk with a Continuum of Random Variables. *Econometrica*, 63(5):1195–1224.

Al-Najjar, N. I. (2004). Aggregation and the law of large numbers in large economies. *Games and Economic Behavior*, 47(1):1–35.

Al-Najjar, N. I. (2008). Large games and the law of large numbers. *Games and Economic Behavior*, 64(1):1–34.

Athey, S. (2002). Monotone Comparative Statics under Uncertainty. *The Quarterly Journal of Economics*, 117(1):187–223.

Aumann, R. and Brandenburger, A. (1995). Epistemic Conditions for Nash Equilibrium. *Econometrica*, 63(5):1161–1180.

Azrieli, Y. (2009). Categorizing others in a large game. *Games and Economic Behavior*, 67(2):351–362.

Börgers, T. and Sarin, R. (1997). Learning Through Reinforcement and Replicator Dynamics. *Journal of Economic Theory*, 77(1):1–14.

Brandenburger, A. (1992). Knowledge and equilibrium in games. *The Journal of Economic Perspectives*, pages 83–101.

Brandenburger, A. and Dekel, E. (1987). Rationalizability and Correlated Equilibria. *Econometrica*, 55(6):1391.

Camerer, C. F. (1997). Progress in Behavioral Game Theory. *The Journal of Economic Perspectives*, 11(4):167–188.

Camerer, C. F., Ho, T.-H., and Chong, J.-K. (2004). A Cognitive Hierarchy Model of Games. *The Quarterly Journal of Economics*, 119(3):861–898.

Choi, M. and Smith, L. (2014). Ordinal Aggregation Results via Karlin's Variation Diminishing Property. *Available at SSRN 2491021*.

Crawford, V. P. (2013). Boundedly Rational versus Optimization-Based Models of Strategic Thinking and Learning in Games. *Journal of Economic Literature*, 51(2):512–527.

Crawford, V. P., Costa-Gomes, M. A., and Iriberri, N. (2013). Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications. *Journal of Economic Literature*, 51(1):5–62.

Eyster, E. and Rabin, M. (2005). Cursed Equilibrium. *Econometrica*, 73(5):1623–1672.

Farouki, R. T. (2012). The Bernstein polynomial basis: A centennial retrospective. *Computer Aided Geometric Design*, 29(6):379–419.

Gavrea, I., Gonska, H. H., and Kacsó, D. P. (1998). On the variation-diminishing property. *Results in Mathematics*, 33(1-2):96–105.

Geanakoplos, J., Pearce, D., and Stacchetti, E. (1989). Psychological games and sequential rationality. *Games and Economic Behavior*, 1(1):60–79.

Haldane, J. B. S. (1932). A note on inverse probability. *Mathematical Proceedings of the Cambridge Philosophical Society*, 28(01):55–61.

Harstad, R. M. and Selten, R. (2013). Bounded-Rationality Models: Tasks to Become Intellectually Competitive. *Journal of Economic Literature*, 51(2):496–511.

Jewitt, I. (1988). Justifying the First-Order Approach to Principal-Agent Problems. *Econometrica*, 56(5):1177.

Jewitt, I. (1991). Applications of Likelihood Ratio Orderings in Economics. *Lecture Notes-Monograph Series*, 19:174–189.

Judd, K. L. (1985). The law of large numbers with a continuum of IID random variables. *Journal of Economic Theory*, 35(1):19–25.

Karlin, S. (1968). *Total Positivity*. Stanford University Press.

Karlin, S. and Ziegler, Z. (1975). Some inequalities for generalized concave functions. *Journal of Approximation Theory*, 13(3):276–293.

Levin, D. and Smith, J. L. (1994). Equilibrium in Auctions with Entry. *The American Economic Review*, 84(3):585–599.

Makowski, A. M. and Philips, T. K. (1992). Stochastic Convexity of Sums of I.I.D. Non-Negative Random Variables with Applications. *Journal of Applied Probability*, 29(1):156.

McKelvey, R. D. and Palfrey, T. R. (1996). A Statistical Theory of Equilibrium in Games*. *Japanese Economic Review*, 47(2):186–209.

Mortensen, D. T. and Pissarides, C. A. (1994). Job Creation and Job Destruction in the Theory of Unemployment. *The Review of Economic Studies*, 61(3):397–415.

Muth, J. F. (1961). Rational Expectations and the Theory of Price Movements. *Econometrica*, 29(3):315–335.

Osborne, M. J. and Rubinstein, A. (1998). Games with Procedurally Rational Players. *The American Economic Review*, 88(4):834–847.

Osborne, M. J. and Rubinstein, A. (2003). Sampling equilibrium, with an application to strategic voting. *Games and Economic Behavior*, 45(2):434–441.

Pallini, A. (2005). Bernstein-type approximations of smooth functions. *Statistica*, 65(2):169–191.

Phillips, G. M. (2003). Bernstein Polynomials. *Interpolation and Approximation by Polynomials*, pages 247–290.

Quah, J. K. H. (2007). The Comparative Statics of Constrained Optimization Problems. *Econometrica*, 75(2):401–431.

Quah, J. K. H. and Strulovici, B. (2012). Aggregating the Single Crossing Property. *Econometrica*, 80(5):2333–2348.

Rabin, M. (2013). Incorporating Limited Rationality into Economics. *Journal of Economic Literature*, 51(2):528–543.

Reny, P. J. (1993). Common Belief and the Theory of Games with Perfect Information. *Journal of Economic Theory*, 59(2):257–274.

Rogers, B. W., Palfrey, T. R., and Camerer, C. F. (2009). Heterogeneous quantal response equilibrium and cognitive hierarchies. *Journal of Economic Theory*, 144(4):1440–1467.

Rowthorn, R. and Sethi, R. (2008). Procedural Rationality and Equilibrium Trust. *The Economic Journal*, 118(530):889–905.

Rubinstein, A. (1991). Comments on the Interpretation of Game Theory. *Econometrica*, 59(4):909–924.

Rubinstein, A. and Wolinsky, A. (1994). Rationalizable Conjectural Equilibrium: Between Nash and Rationalizability. *Games and Economic Behavior*, 6(2):299–311.

Rustichini, A. (2003). Equilibria in large games with continuous procedures. *Journal of Economic Theory*, 111(2):151–171.

Schelling, T. C. (2006). *Micromotives and Macrobehavior*. W. W. Norton & Company, New York, Revised edition edition.

Selten, R. and Chmura, T. (2008). Stationary Concepts for Experimental 2x2-Games. *The American Economic Review*, 98(3):938–966.

Sethi, R. (2000). Stability of Equilibria in Games with Procedurally Rational Players. *Games and Economic Behavior*, 32(1):85–104.

Shaked, M. (1980). On Mixtures from Exponential Families. *Journal of the Royal Statistical Society. Series B (Methodological)*, 42(2):192–198.

Shaked, M. (1982). Dispersive Ordering of Distributions. *Journal of Applied Probability*, 19(2):310–320.

Simon, H. A. (1955). A Behavioral Model of Rational Choice. *The Quarterly Journal of Economics*, 69(1):99–118.

Stancu, D. D. (1979). Application of divided differences to the study of monotonicity of the derivatives of the sequence of Bernstein polynomials. *Calcolo*, 16(4):431–445.

Yildiz, M. (2007). Wishful Thinking in Strategic Environments. *The Review of Economic Studies*, 74(1):319–344.

Zhu, M. and Lu, A. Y. (2004). The counter-intuitive non-informative prior for the Bernoulli family. *Journal of Statistics Education*, 12(2):1–10.