

Quantal Response Equilibria with Heterogeneous Agents

Russell Golman^{a*}

April 24, 2011

^a Department of Applied and Interdisciplinary Mathematics, University of Michigan,
2082 East Hall, 530 Church Street, Ann Arbor, MI 48109, USA. Telephone: 734 764-
7436. Fax: 734 763-0937. E-mail: rgolman@umich.edu

*I am grateful to Andreas Blass and Scott Page for many helpful conversations.

Abstract

This paper introduces a model of quantal response equilibrium with heterogeneous agents and demonstrates the existence of a representative agent for such populations. Except in rare cases, the representative agent's noise terms cannot be independently and identically distributed across the set of actions, even if that is assumed for the individual agents. This result demonstrates a fundamental difference between a representative agent and truly heterogeneous quantal responders and suggests that when fitting quantal response specifications to aggregate data from a population of subjects, the noise terms should be allowed to be jointly dependent across actions. Even though this introduces additional degrees of freedom, it makes the model well specified. The representative agent inherits a regular quantal response function from the actual agents, so this model does impose falsifiable restrictions.

KEYWORDS: quantal response equilibria, bounded rationality, representative agent, heterogeneity

JEL classification codes: C72, C50

1 Introduction

Quantal response equilibrium extends the Nash Equilibrium notion to allow bounded rationality. Players can be seen as making errors while trying to choose optimal strategies, or equivalently, as observing payoffs disturbed by idiosyncratic noise. Players may select any action with positive probability, as assigned by their quantal response functions.

Here we examine the use of a representative agent to describe the aggregate behavior of heterogeneous quantal responders. Different agents may have different distributions of payoff noise and hence different quantal response functions. A representative agent would have the average quantal response function. Weakening a standard assumption about the admissible distributions of payoff noise, we show existence of a representative agent. However, this representative agent does not have a representative distribution of payoff noise, nor any iid distribution in large enough games. Nevertheless, important regularity properties are inherited by a representative agent, and thus quantal response equilibrium models make testable population-wide predictions for heterogeneous populations.

We consider structural quantal response equilibria (QRE) [17, 4] in the context of a population game. In a large population of agents, we should expect heterogeneity of behavior [12, 15]. A population of quantal responders should consist of agents who may have different error rates, or different distributions of payoff noise. In fact, McKelvey, et. al. [18] find experimental evidence for heterogeneous error distributions

in trying to fit logit QRE to data on two-by-two asymmetric games. Weizsäcker [23] also finds evidence for heterogeneous error distributions in a more general context.¹

We allow for arbitrary distributions of payoff noise. Prior research into quantal response equilibria with heterogeneous agents has considered a distribution of parameters which parametrize the distributions of payoff noise [21, 23], with particular interest in distributions of logit responders [20]. Here, we model heterogeneous distributions of payoff noise with a functional defined over distribution functions. As we do not assume that distributions of payoff noise take any particular functional forms, this approach allows for more distribution functions than can be described with finitely many parameters.

Our interest is in the behavior of an entire population, and we seek a representative agent whose mixed strategy quantal response always matches the population aggregate. A representative agent would behave like the population as a whole in any normal-form game. We need representative-agent models because while we believe people really are heterogeneous, in practice it would be quite tedious to determine each person's quantal response function individually when we fit data. To achieve statistically significant sample sizes, data is typically pooled across subjects. The representative agent is what we can estimate in an experiment. Thus, a representative-agent model enables us to capture aggregate behavior while recognizing underlying

¹Further motivation to consider heterogeneity in a population of quantal responders comes from recent findings that models of heterogeneous learners often cannot be adequately approximated by representative-agent models with common parameter values for all [24, 8, 6].

heterogeneity.

With weak assumptions on the agents' distributions of payoff noise we prove existence of a representative agent. However, the distribution of payoff disturbances necessary to produce representative choices is not representative of the noise the actual agents observe in their payoffs. We show that in games with enough pure strategies, a representative agent could not have payoff disturbances independent and identically distributed across actions even if the actual agents did. Does this mean a representative-agent QRE model can rationalize any aggregate choice behavior, even when agents in the population have payoff disturbances i.i.d. across actions [7]? No! Quantal response equilibrium with representative agents is a concept not void of empirical content. We find that if agents all use regular quantal response functions (as defined by Goeree, et. al. [4]), then the representative agent's quantal response must also be regular. Falsifiable predictions are generated by the regularity of the representative agent's quantal response function, not by the specification of its payoff noise.

The most prominent form of quantal response is the logit specification. Notably, the representative agent for a population of heterogeneous logit responders is not another logit responder. A companion paper considers heterogeneous logit responders and finds that a mis-specified homogeneous logit parameter exhibits a systematic downward bias [5]. This implies that empirical estimates disregarding heterogeneity mis-specify agents as *less rational* than they really are, at least on average. Such a mis-specified model lacks predictive power across games. A representative-agent

model, on the other hand, would accurately describe behavior in any game.

Population games capture anonymous interaction with large numbers of participants for all roles. Suppose players in a game do not know the identity of the other players, but only have information about the populations of potential players. They may treat the population-wide distribution of opponents' choices as arising from a representative agent playing a mixed strategy. In modeling collective play of the game, we may do the same.

Nash proposed a population game interpretation of equilibrium in his unpublished PhD dissertation [22]. Following his lead, we assume that there is a population of agents for each role in a game. A generic n -player game involves n populations of agents, but if multiple players have identical roles and we adopt the restriction that players in identical roles should play identical population mixed strategies, then these players may be selected from the same population. So, in a totally symmetric game, we may have only a single population of agents. Similarly, in a discrete choice environment, a single player would be a randomly selected member of a single population. We assume the populations are large, and we are interested in the fraction of a population playing a given strategy. An agent's payoff in a population game is the average of his payoffs against all other combinations of agents (or equivalently his expected payoff given random matching).

Population games provide a framework for the use of evolutionary learning dynamics. Learning rules that assume that players noisily best respond often converge to QRE [3, 13, 10, 14, 2, 9, 11]. This paper focuses on the QRE itself and not on

any particular learning rule that might lead to it. Population games also describe experimental settings well (at least when anonymity is maintained in the lab), as data is accumulated through the randomly matched interactions of many subjects.

This paper is organized as follows. Section 2 introduces the notation in the context of a single population and provides definitions of a QRE and a representative agent. Section 3 contains our general results describing a representative agent. In Section 4, we extend our framework and our results to n -player asymmetric games. Section 5 concludes. The Appendix contains two technical lemmas.

2 A Single Population

To simplify the presentation, we begin with a single population of agents. The context can be thought of as a symmetric game or alternatively a one-player game, i.e., a discrete choice environment with additive random utility. In Section 4, we see that our results apply to general n -player asymmetric games as well. For mathematical generality we allow the population to be uncountable, but we make no assumption it is so. We require only that the population is sufficiently large so that a single agent has negligible impact on collective play.

Let $S = \{s_1, \dots, s_J\}$ be the set of pure strategies available to the agents. The collective play of all the agents defines the population mixed strategy x . Formally, $x \in \Delta^{J-1}$, the $(J-1)$ -dimensional simplex where $x_j \geq 0$ for all j and $\sum_j x_j = 1$. We interpret x_j as the fraction of the population selecting pure strategy s_j .

2.1 Payoff Noise

A structural QRE arises when agents' utility functions are modified by noise terms, privately observed stochastic payoff disturbances. Denote by π_j the payoff from taking pure strategy s_j . Of course, payoffs are a function of the strategies used by all the players, $\pi_j = \pi_j(x)$, but we omit the function's argument for ease of notation. Also denote the vector $\pi = \pi_1, \dots, \pi_J$. Formally, $\pi : \Delta^{J-1} \rightarrow \Re^J$. For each pure strategy s_j , agent μ observes a payoff disturbance ϵ_j^μ , making agent μ 's disturbed payoff $\pi_j^\mu = \pi_j + \epsilon_j^\mu$. This is the function agents maximize with their choice of strategy in a QRE.

The distribution of payoff disturbances is assumed to be *admissible*, meaning that:

- (a1) the disturbances are independent across agents;
- (a2) each agent has an absolutely continuous joint distribution of $(\epsilon_1^\mu, \dots, \epsilon_J^\mu)$ that is independent of the population mixed strategy x , i.e., all marginal densities exist;
- (a3) disturbances are unbiased in the sense that they all have mean zero.

Allowing only admissible distributions guarantees the existence of a QRE [17]. Here, we make the additional assumption that for each agent, disturbances are independent and identically distributed (iid) across the set of actions. This assumption could be relaxed, but some such restriction is necessary for the QRE notion to produce falsifiable predictions [7].

When the setup for QRE does not explicitly involve populations of agents, it is assumed that each player has a distribution of payoff disturbances. In the context of a population game, this corresponds to each agent within the population having an identical distribution of disturbances. That is, the convention is to assume homogeneous populations.² Here, we specifically want to leave open the possibility that agents in the same population have different distributions of payoff shocks. So, we do *not* assume identical distributions of ϵ_j^μ for all μ .

2.2 Heterogeneity

To model heterogeneity in the distributions of payoff disturbances, we will construct a “probability (mass) density functional” defined over such distributions. Let $P_\epsilon(\cdot)$ be a distribution function for the payoff disturbance to a particular action. Each agent has a distinct P_ϵ , which is the same for all actions in that agent’s strategy space. Let \mathcal{P} denote the space of all admissible distribution functions for an arbitrary payoff disturbance, i.e., all unbiased, absolutely continuous distribution functions on \mathfrak{R} . Formally, we capture heterogeneity by defining a probability measure G_ϵ on the space \mathcal{P} . To make our presentation more accessible, we will reformulate it with a functional $F_\epsilon[P_\epsilon]$ that associates to each distribution function P_ϵ a probability mass or

²In the discrete choice literature, the convention is to assume heterogeneous agents who can perfectly observe their utilities [16, 1]. The formal representation is the same as for homogeneous agents who observe disturbed utilities. In this paper we want to consider heterogeneous agents who observe disturbed utilities.

density describing the fraction of the population with payoff disturbances distributed by P_ϵ . Using the Radon-Nikodym theorem, for any arbitrary measure ν on \mathcal{P} such that G_ϵ is absolutely continuous with respect to ν , there exists a functional F_ϵ (called the Radon-Nikodym derivative of G_ϵ) such that $G_\epsilon(\mathbf{p}) = \int_{P_\epsilon \in \mathbf{p}} F_\epsilon[P_\epsilon] \nu(dP_\epsilon)$ for any measurable $\mathbf{p} \subseteq \mathcal{P}$. In an abuse of notation justified by our desire not to worry about the particular form of the measure ν , we will use $\int dP_\epsilon$ as a shorthand for the Lebesgue integral $\int \nu(dP_\epsilon)$. This notation will be used even when there are mass points and is not meant to suggest their exclusion. We also adopt the convention of integrating over the entire space unless otherwise specified.

In this approach the functional F_ϵ captures a distribution of distributions of payoff shocks in the population. It thus provides a general way to think about heterogeneity of quantal responses. The conventional assumption of a homogeneous population can be recaptured, for example, by taking $F_\epsilon[P_\epsilon] = 1$ for a particular P_ϵ and 0 everywhere else. Alternatively, a heterogeneous logit response model [20, 5] arises by taking the support of F_ϵ to be (type I) extreme value distributions with a range of parameter values.

Strategy choices are determined by comparing the disturbed payoffs, so the most relevant variables are the differences between payoff shocks, $\delta_{jj'}^\mu = \epsilon_j^\mu - \epsilon_{j'}^\mu$. These $\delta_{jj'}^\mu$ are identically distributed across all j and $j' \neq j$ because the ϵ_j^μ are iid across all j . By absolute continuity, the marginal densities exist, and they are even functions because the $\delta_{jj'}^\mu$ are antisymmetric in the indices j and j' . There is obviously dependence among these random variables across j and j' . We will consider the

$(J - 1)$ -dimensional random vector $\delta_j^\mu = \left(\delta_{1j}^\mu, \dots, \widehat{\delta_{jj}^\mu}, \dots, \delta_{Jj}^\mu \right)$ for an arbitrary j , which then determines the value of all other δ_j^μ . The hat notation indicates that an element is removed from the vector – in this case δ_{jj}^μ is removed because by definition it is always 0. Note that admissibility of the payoff disturbances implies δ_j^μ has zero mean because all the ϵ_j^μ have zero mean.

Heterogeneity in the distributions of payoff shocks leads to heterogeneity in the distributions of the differences between payoff shocks. That is, the functional over P_ϵ induces a functional defined on the joint distributions for δ_j^μ . Let $P : \mathfrak{R}^{J-1} \rightarrow [0, 1]$ be a joint distribution function of δ_j^μ for some μ and any $j \in [1, J]$. We construct the functional $F[P]$ as follows. Let $\mathcal{P}_{\epsilon \rightarrow \delta}(P)$ be the set of $P_\epsilon(\cdot)$, distribution functions of ϵ_j^μ , which give rise to $P(\cdot)$. Let \mathcal{P}_δ be the space of all $(J - 1)$ -dimensional joint distribution functions. Taking measures $\nu_{\epsilon \rightarrow \delta}$ and ν_δ on $\mathcal{P}_{\epsilon \rightarrow \delta}(P)$ and \mathcal{P}_δ respectively, such that $\nu(\mathbf{p}) = \int_{P_\epsilon \in \mathcal{P}_{\epsilon \rightarrow \delta}(P)} \nu_{\epsilon \rightarrow \delta}(\mathbf{p} \cap \mathcal{P}_{\epsilon \rightarrow \delta}(P)) \nu_\delta(dP)$, we arrive at

$$F[P] = \int_{P_\epsilon \in \mathcal{P}_{\epsilon \rightarrow \delta}(P)} F_\epsilon[P_\epsilon] \nu_{\epsilon \rightarrow \delta}(dP_\epsilon).$$

Note that there are joint distribution functions that could not apply to any δ_j^μ because they do not describe differences between iid random variables, and our definition implies $F = 0$ for these functions.

2.3 Quantal Response

The quantal response function for each agent returns the agent's likelihood of choosing each strategy given the agent's undisturbed payoffs. Let $Q_j^\mu(\pi)$ be the probability that

agent μ selects strategy s_j given the payoffs to each strategy. Formally, for any vector $\pi' = (\pi'_1, \dots, \pi'_J) \in \mathfrak{R}^J$, define

$$R_j^\mu(\pi') = \{(\varepsilon_1^\mu, \dots, \varepsilon_J^\mu) \in \mathfrak{R}^J : \pi_j + \varepsilon_j^\mu \geq \pi_{j'} + \varepsilon_{j'}^\mu \text{ for all } j' = 1, \dots, J\}$$

to be the set of realizations of agent μ 's joint set of payoff disturbances that would lead to choosing action s_j . Then $Q_j^\mu(\pi) = \text{Prob}\{(\varepsilon_1^\mu, \dots, \varepsilon_J^\mu) \in R_j^\mu(\pi)\}$. Using the joint distribution function for the $(J-1)$ -dimensional random vector δ_j^μ , the j^{th} component of the quantal response function can be expressed as

$$Q_j^\mu(\pi) = P(\pi_j - \pi_1, \pi_j - \pi_2, \dots, \pi_j - \pi_J), \quad (1)$$

naturally omitting $\pi_j - \pi_j$ just as we did δ_{jj}^μ . Thus, an agent's quantal response function is determined by the joint distribution of differences between payoff shocks.

The quantal response functions for all the agents can be aggregated across the population to give the population mixed strategy response to any given population state. In a finite population of m agents, the population aggregate quantal response is $Q_j = \frac{1}{m} \sum_{\mu=1}^m Q_j^\mu$ for all j . More generally, the aggregate quantal response in an infinite population is

$$Q_j = \int F_\varepsilon[P_\varepsilon] Q_j^\mu dP_\varepsilon \quad (2)$$

where we abuse notation by letting $\mu = \mu(P_\varepsilon)$ be an agent with payoff disturbances iid from P_ε . This is just the expectation of agents' quantal response functions with respect to the probability mass / density functional F_ε . It can be taken pointwise, i.e., independently for every value of the payoff vector π .

We can now define a quantal response equilibrium and then formally describe a representative agent for this heterogeneous population.

Definition A *quantal response equilibrium (QRE)* is defined by the fixed point equation $x_j = Q_j(\pi(x))$ for all j .

Whereas a Nash Equilibrium is a state of play with everybody simultaneously playing a best response, a QRE is a state with everybody simultaneously playing according to their quantal response functions.

Definition A *representative agent* would have a quantal response function $\hat{Q}(\pi)$ equal to the population aggregate quantal response function:

$$\hat{Q} = (Q_1, \dots, Q_J). \quad (3)$$

For all games, the population as a whole behaves exactly as if it were homogeneously composed of representative agents.

Our definition of a representative agent can now be translated into a statement about the representative joint distribution of differences between payoff shocks. It means the representative agent would have δ_j distributed according to a joint distribution function $\hat{P}(\cdot)$ such that

$$\hat{P} = \int F[P] P dP. \quad (4)$$

We can think of \hat{P} as the population's expected joint distribution function for differences between payoff shocks, with respect to the induced probability mass / density

functional F .³ The representative agent's quantal response function can in turn be found by using \hat{P} in Eq. (1). This provides a working definition of a representative agent that is more useful than Eq. (3).

Example We demonstrate finding the representative agent's quantal response function with $J = 2$ in this simple example. Suppose the population consists of agents having payoff disturbances iid from an exponential distribution (shifted to have mean zero) with parameters ranging from λ_{\min} to λ_{\max} according to a uniform distribution. This population is described by the functional $F_\epsilon[P_\epsilon]$ with support on distribution functions such that $P_\epsilon(\epsilon) = 1 - e^{-\lambda(\epsilon + \frac{1}{\lambda})}$ over $\epsilon \in [-\frac{1}{\lambda}, \infty)$, for $\lambda_{\min} \leq \lambda \leq \lambda_{\max}$; and with constant value $\frac{1}{\lambda_{\max} - \lambda_{\min}}$ on this support. The difference between two iid *Exponential*(λ) random variables has a *Laplace*($0, \frac{1}{\lambda}$) distribution. Thus, the induced functional $F[P]$ has support on distribution functions such that

$$P(\delta) = \begin{cases} \frac{1}{2}e^{\lambda\delta} & \text{if } \delta \leq 0 \\ 1 - \frac{1}{2}e^{-\lambda\delta} & \text{if } \delta > 0 \end{cases}$$

for $\lambda_{\min} \leq \lambda \leq \lambda_{\max}$. Integrating over the uniform distribution of types gives us

$$\hat{P}(\delta) = \begin{cases} \frac{1}{2\delta(\lambda_{\max} - \lambda_{\min})} (e^{\lambda_{\max}\delta} - e^{\lambda_{\min}\delta}) & \text{if } \delta < 0 \\ \frac{1}{2} & \text{if } \delta = 0 \\ 1 + \frac{1}{2\delta(\lambda_{\max} - \lambda_{\min})} (e^{-\lambda_{\max}\delta} - e^{-\lambda_{\min}\delta}) & \text{if } \delta > 0. \end{cases}$$

Using Eq. (1), we then obtain the representative agent's quantal response function.

If $\pi_1 = \pi_2$, then $\hat{Q}(\pi) = (\frac{1}{2}, \frac{1}{2})$. Otherwise, assuming without loss of generality that

³We again abuse notation by neglecting to specify the measure ν_δ .

$\pi_1 < \pi_2$, we have

$$\hat{Q}(\pi) = \left(\frac{e^{\lambda_{\max}(\pi_1 - \pi_2)} - e^{\lambda_{\min}(\pi_1 - \pi_2)}}{2(\pi_1 - \pi_2)(\lambda_{\max} - \lambda_{\min})}, 1 - \frac{e^{\lambda_{\max}(\pi_1 - \pi_2)} - e^{\lambda_{\min}(\pi_1 - \pi_2)}}{2(\pi_1 - \pi_2)(\lambda_{\max} - \lambda_{\min})} \right).$$

2.4 Characteristic Functions

The property of having the population aggregate quantal response function was our first and most transparent characterization of a representative agent. We then established that this is equivalent to having the representative joint distribution function for the vector of differences between payoff shocks as described in Eq. (4). We now develop another equivalent characterization that will soon prove useful.

Given a functional that describes the heterogeneity of the population, we can use characteristic functions to identify a representative agent. This approach is effective because there is a bijection between distribution functions and characteristic functions. Let $\theta : \Re \rightarrow \mathbb{C}$ be the characteristic function of a payoff disturbance ϵ_j with distribution function $P_\epsilon(\cdot)$,

$$\theta(t) = E(e^{it\epsilon_j}).$$

Note that θ is a complex-valued function of a single real variable and has complex conjugate $\bar{\theta}(t) = \theta(-t)$. It must be uniformly continuous, satisfy $\theta(0) = 1$ and $|\theta(t)| \leq 1$, and the quadratic form in u and v with $\theta(u - v)$ as its kernel must be non-negative definite. (These properties can be used to define an arbitrary characteristic function.) Take $\phi : \Re^{J-1} \rightarrow \mathbb{C}$ to be the characteristic function associated with the joint distribution $P(\cdot)$ of δ_j . We still write $\phi(t)$, now assuming $t = (t_1, \dots, t_{J-1})$ to

be a vector in \Re^{J-1} . We can express ϕ in terms of θ ,

$$\begin{aligned}
\phi(t) &= E(e^{it \cdot \delta_j}) \\
&= E(e^{it_1 \delta_{1j}} \dots e^{it_{J-1} \delta_{Jj}}) \\
&= E(e^{it_1(\epsilon_1 - \epsilon_j)} \dots e^{it_{J-1}(\epsilon_{J-1} - \epsilon_j)}) \\
&= E(e^{it_1 \epsilon_1}) \dots E(e^{it_{J-1} \epsilon_{J-1}}) \cdot E(e^{-i(\sum_{l=1}^{J-1} t_l) \epsilon_j}) \\
&= \theta(t_1) \dots \theta(t_{J-1}) \cdot \theta(-\sum_{l=1}^{J-1} t_l). \tag{5}
\end{aligned}$$

In addition to the properties just mentioned, we also know that if $\sum_{l=1}^J r_l = 0$, then $\phi(r_1, \dots, \widehat{r_j}, \dots, r_J)$ is independent of j , because by Eq. (5) it has the same expansion in terms of θ for all j . If there are only two actions, $J = 2$, then ϕ is real and positive because P is symmetric (see Eq. (5), where $\phi(t)$ would be the product of $\theta(t)$ and its complex conjugate). The functional F_ϵ induces a distribution over characteristic functions $\Psi_\epsilon[\theta] = F_\epsilon[P_\epsilon]$. Similarly, define $\Psi[\phi] = F[P]$.

Let

$$\widehat{\phi}(t) = \int \Psi[\phi] \phi(t) d\phi \tag{6}$$

be the expectation of characteristic functions for δ_j^μ in the population.⁴ This representative characteristic function can be constructed by taking the integral pointwise, i.e., independently for every value of t . Fixing the input point t , we know that this functional integral always converges because $|\phi(t)| \leq 1$.

⁴We repeat our earlier abuse of notation, again neglecting to specify the integration measure.

3 A Representative Agent

3.1 Existence of a Representative Agent

The first issue to address is whether a representative agent always exists. Theorem 1 tells us that there is only one pathological type of heterogeneity for which the population does not have a representative agent, and Corollaries 1 and 2 tell us not to get hung up on the technicality that causes this problem in an infinite population, as it is of no practical importance. The representative joint distribution function $\hat{P}(\cdot)$ can always be constructed given the functional $F[P]$ describing the heterogeneity in the population, but there is a danger that it is not an admissible distribution function. Specifically, it may fail to have finite mean. This may occur only when infinitely many different distribution functions are in use in the population. As a corollary to the theorem, a representative agent is sure to exist in all finite-type models. Additionally, relaxing the requirement that distributions of disturbances must have zero mean with a more expansive definition of admissible unbiased payoff noise also ensures the existence of a representative agent.

Theorem 1 *Define $\hat{P}(\cdot)$ as in Eq. (4).⁵ If $\hat{P}(\cdot)$ has finite mean, then a representative agent exists with δ_j distributed by $\hat{P}(\cdot)$ and having characteristic function $\hat{\phi}(t)$.*

⁵Recall that the integral in Eq. (4) is a functional integral over possible distribution functions for payoff noise, not an integral with respect to the probability measure of a particular noise distribution. This integral is used point-by-point to define $\hat{P}(\cdot)$, the representative agent's joint distribution function for differences between payoff shocks. Existence of the mean of $\hat{P}(\cdot)$ then depends on an integral with respect to this particular distribution, and is not the same as existence of the functional

Proof It is well known that the maps between distribution functions and characteristic functions are linear. Apply the Levy continuity theorem to Eq. (6). This requires $\hat{\phi}(t)$ to be continuous at $t = 0$, which we establish with Lemma 1 in the Appendix. Lemma 2 in the Appendix establishes that the mean of $\hat{P}(\cdot)$ is 0 if it exists, and thus $\hat{P}(\cdot)$ is admissible when this is the case. ■

Corollary 1 *If $F[P] > 0$ for only finitely many joint distribution functions $P(\cdot)$, then a representative agent exists.*

Proof The only source for divergence of the mean of $\hat{P}(\cdot)$ is the limit that results from $F[P] > 0$ for infinitely many P . All the joint distribution functions in the support of F have zero mean, so a finite linear combination of them also describes a random vector with zero mean. Then Theorem 1 applies. ■

Taking a closer look at an example $\hat{P}(\cdot)$ that has divergent mean and thus fails to be an admissible joint distribution function offers insight into how such cases arise.

Example For simplicity, assume $J = 2$. The example works just as well with more pure strategies, but the notation becomes cluttered. Partition the set of joint distribution functions $P(\cdot)$ into \mathcal{P}_y such that $P(\cdot) \in \mathcal{P}_y$ implies $P(e^y) \leq 1 - \alpha$ for some fixed positive $\alpha < \frac{1}{2}$. This partition is not uniquely determined, but as long as the

integral that defines $\hat{P}(\cdot)$. The admissibility condition requiring existence of this mean is simply a convention.

\mathcal{P}_y are non-empty, it will do. Consider the functional $F[P]$ where

$$\int_{\mathcal{P}_y} F[P] dP = \begin{cases} e^{-y} & \text{for } y \geq 0 \\ 0 & \text{for } y < 0. \end{cases}$$

Then the mean of $\hat{P}(\cdot)$ is divergent because

$$\begin{aligned} \int_0^\infty \delta d\hat{P}(\delta) &= \int_0^\infty \delta \int_0^\infty \int_{\mathcal{P}_y} F[P] P'(\delta) dP dy d\delta \\ &\geq \int_0^\infty \int_{\mathcal{P}_y} F[P] \int_{e^y}^\infty \delta P'(\delta) d\delta dP dy \\ &\geq \int_0^\infty \int_{\mathcal{P}_y} F[P] \alpha e^y dP dy \\ &= \int_0^\infty \alpha dy. \end{aligned}$$

Admissibility requires $\hat{P}(\cdot)$ to have zero mean, but when this fails, we shouldn't conclude that a representative quantal response function does not exist. Instead, we can relax the requirements of admissibility to guarantee that a representative agent always exists. The restriction to zero mean payoff disturbances is not necessary for the existence of a QRE, as fixed point theorems can be applied without it. The desire for unbiased disturbances appears to be aesthetic, and the possible inadmissibility of representative agents is an artifact of the way it is implemented. Consider replacing the zero mean assumption (a3) with the following alternative:

(a3') the Cauchy principal value of the mean of each payoff disturbance is zero⁶, and

$$\lim_{\gamma \rightarrow \infty} \gamma \text{Prob} \{ |\epsilon_j^\mu| \geq \gamma \} = 0 \text{ for each } \epsilon_j^\mu.$$

⁶The Cauchy principal value of an improper integral $\int_{-\infty}^\infty f(t) dt$ is defined as $\lim_{T \rightarrow \infty} \int_{-T}^T f(t) dt$.

Assumption (a3') holds whenever assumption (a3) is satisfied, so this is a weaker condition to impose on the payoff disturbances. Even though the mean of ϵ_j^μ may blow up under assumption (a3'), these disturbances are still unbiased, and their likelihood still decays sufficiently quickly as they get large.

Definition We say payoff disturbances are *weakly admissible* if assumptions (a1), (a2) and (a3') hold.

With just this slight relaxation of admissibility, we always get a representative agent.

Corollary 2 *Allow weakly admissible payoff disturbances. A representative agent exists with δ_j distributed by $\hat{P}(\cdot)$ and having characteristic function $\hat{\phi}(t)$.*

Proof Lemma 2 shows that $\hat{P}(\cdot)$ always satisfies the weak admissibility assumption (a3'). In turn, there exists a joint distribution of $(\epsilon_1, \dots, \epsilon_J)$ that satisfies (a3') and is consistent with δ_j being distributed by $\hat{P}(\cdot)$. ■

3.2 Payoff Disturbances for the Representative Agent

We have defined a representative agent with the property that the agent's choice of strategy is representative of the population as a whole. We now show that this is not equivalent to having representative noise in the underlying payoffs. The mapping from payoff disturbances to choices is nonlinear, and we should not expect averaging over all agents to preserve this map. We say $\hat{P}_\epsilon(\cdot)$ is a *representative distribution of payoff shocks* if it is a (weakly) admissible distribution function such that

$$\hat{P}_\epsilon = \int F_\epsilon[P_\epsilon] P_\epsilon dP_\epsilon. \tag{7}$$

By applying the Levy continuity theorem here too, we find that a representative distribution of payoff shocks has characteristic function $\hat{\theta}(t) = \int \Psi_\epsilon[\theta] \theta(t) d\theta$. With this groundwork in place, we are ready for Theorem 2, which says that a representative quantal response function does not arise from a representative distribution of payoff shocks.

Theorem 2 *A representative agent has a representative distribution of payoff shocks if and only if the population is homogeneous.*

Proof Let Θ be the set of characteristic functions of ϵ_j that give rise to a given $\phi(\cdot)$. Using Eq. (5), $\Theta = \{\theta : \phi(t) = \left(\prod_{l=1}^{J-1} \theta(t_l)\right) \cdot \theta(-\sum_{l=1}^{J-1} t_l)\}$. From the relationships between the functionals, we have $\Psi[\phi] = \int_{\Theta} \Psi_\epsilon[\theta] d\theta$. We can then express a representative agent's characteristic function for δ_j as

$$\hat{\phi}(t) = \int \Psi_\epsilon[\theta] \left(\prod_{l=1}^{J-1} \theta(t_l)\right) \theta(-\sum_{l=1}^{J-1} t_l) d\theta.$$

But

$$\int \Psi_\epsilon[\theta] \left(\prod_{l=1}^{J-1} \theta(t_l)\right) \theta(-\sum_{l=1}^{J-1} t_l) d\theta \neq \left(\prod_{l=1}^{J-1} \int \Psi_\epsilon[\theta] \theta(t_l) d\theta\right) \cdot \int \Psi_\epsilon[\theta] \theta(-\sum_{l=1}^{J-1} t_l) d\theta \quad (8)$$

unless for each t_l , $\theta(t_l)$ is the same for all θ in the support of Ψ_ϵ . Since t_l is an arbitrary variable, this would mean there could only be one function in the support of Ψ_ϵ , i.e., no heterogeneity of distributions of payoff shocks in the population. ■

In light of the fact that a representative agent for a heterogeneous population does not have a representative distribution of payoff shocks, the question arises as to what

distribution of payoff shocks could actually produce a representative agent. According to the next result, if there are enough actions and there is heterogeneity of the $\delta_{jj'}$, then the representative agent cannot arise from any distribution of payoff shocks that is iid across the set of actions. Theorem 3 says that if there are just two actions, there is an iid distribution of payoff shocks (possibly many such distributions) that generates the representative agent. But, if there are at least four actions, assuming heterogeneity of the $\delta_{jj'}$, it is impossible for an iid distribution of payoff shocks to generate the representative agent.⁷

Theorem 3 *Given a representative agent, if $J = 2$, there exists a distribution of payoff shocks iid across all actions and each with characteristic function $\hat{\vartheta}(\cdot)$ such that*

$$\hat{\phi}(t) = \left(\prod_{l=1}^{J-1} \hat{\vartheta}(t_l) \right) \cdot \hat{\vartheta}\left(-\sum_{l=1}^{J-1} t_l\right). \quad (9)$$

But, when $J \geq 4$, there is no $\hat{\vartheta}(\cdot)$ that satisfies Eq. (9) unless every $P_\epsilon(\cdot)$ in the support of F_ϵ gives the same distribution of the $\delta_{jj'}$.

Proof When $J = 2$, we must find a $\hat{\vartheta}(\cdot)$ such that $\hat{\phi}(t_1) = \hat{\vartheta}(t_1) \cdot \hat{\vartheta}(-t_1)$. Recall $J = 2$ implies that all $\phi(\cdot)$ are real and positive, and hence so is $\hat{\phi}$. It suffices to take $\hat{\vartheta}(t_1) = \hat{\vartheta}(-t_1) = \sqrt{\hat{\phi}(t_1)}$.

⁷Examples indicate that when there are three actions, the representative agent usually cannot arise from iid shocks, but we cannot rule out special cases of heterogeneity for which the representative agent is compatible with iid disturbances. When jointly dependent disturbances are necessary, the correlation structure (implied by Eq. (6)) depends on the makeup of the population.

Now consider $J \geq 4$. Given that individual agents do have payoff shocks that are iid across all actions, any $\phi(\cdot)$ in the population can be expressed in terms of $\theta(\cdot)$ with Eq. (5). Specifically, $\phi(a, -a, a, 0, \dots, 0) = (\theta(a)\theta(-a))^2$. Similarly, $\phi(a, 0, \dots, 0) = \theta(a)\theta(-a)$. Thus,

$$\phi(a, -a, a, 0, \dots, 0) = (\phi(a, 0, \dots, 0))^2.$$

But

$$\int \Psi[\phi] \phi(a, -a, a, 0, \dots, 0) d\phi \neq \left(\int \Psi[\phi] \phi(a, 0, \dots, 0) d\phi \right)^2$$

unless there is no variance of $\theta(a)\theta(-a)$ in the population. Note that $\delta_{jj'}$ has characteristic function $\theta(t)\theta(-t)$. Thus, if there are two distribution functions in the support of $F_\epsilon[P_\epsilon]$ that give different distributions of $\delta_{jj'}$, then for some a ,

$$\hat{\phi}(a, -a, a, 0, \dots, 0) \neq \left(\hat{\phi}(a, 0, \dots, 0) \right)^2.$$

This would mean $\hat{\phi}(\cdot)$ could not be expressed as $\left(\prod_{l=1}^{J-1} \hat{\vartheta}(t_l) \right) \cdot \hat{\vartheta}(-\sum_{l=1}^{J-1} t_l)$ for any $\hat{\vartheta}(\cdot)$. ■

Theorem 3 sounds a cautionary note that even if we believe all agents have noise in their payoffs that is iid across their actions, heterogeneity of the agents leads the population as a whole to behave as if payoff disturbances were not iid across actions.

We desired agents with payoff noise iid across actions because this assumption imposes restrictions on behavior that can be tested empirically. Although it turns out the representative agent may not have payoff noise iid across actions, the representative-agent notion still has empirical content because some properties are inherited from the underlying agents.

3.3 Regularity of a Representative Agent

Goeree, et. al. [4] introduce four axioms which define a *regular quantal response function* $Q^\mu : \mathfrak{R}^J \rightarrow \Delta^{J-1}$ without reference to payoff noise:

(A1) *Interiority*: $Q_j^\mu(\pi) > 0$ for all $j = 1, \dots, J$ and for all $\pi \in \mathfrak{R}^J$.

(A2) *Continuity*: $Q_j^\mu(\pi)$ is a continuous and differentiable function for all $\pi \in \mathfrak{R}^J$.

(A3) *Responsiveness*: $\frac{\partial Q_j^\mu(\pi)}{\partial \pi_j} > 0$ for all $j = 1, \dots, J$ and for all $\pi \in \mathfrak{R}^J$.

(A4) *Monotonicity*: $\pi_j > \pi_{j'}$ implies $Q_j^\mu(\pi) > Q_{j'}^\mu(\pi)$, for all $j, j' = 1, \dots, J$.

They argue that all quantal response functions obey Continuity and weakly obey Responsiveness. If the density of payoff disturbances has full support, then Interiority and Responsiveness are strictly satisfied. When payoff disturbances are iid across actions, then the quantal response function obeys Monotonicity as well. The condition of having payoff disturbances iid across actions is stronger than that of having a monotonic quantal response function in that the former is sufficient, but not necessary, for Monotonicity.

We now show that any regularity property that holds for the underlying agents in the population also holds for the representative agent.

Theorem 4 *If a regularity axiom $\{(A1), (A2), (A3), \text{ or } (A4)\}$ applies to Q^μ for all μ (i.e., for $\mu = \mu(P)$ whenever $P(\cdot)$ is in the support of F), then that axiom applies to the representative agent's quantal response function \hat{Q} .*

Proof Continuity holds for all quantal response functions as a result of the admissibility assumption (a2) that distributions of payoff noise must be absolutely continuous [4]. Interiority, Responsiveness, and Monotonicity each follow from Eq. (2) and (3), which define a representative agent. Essentially, we just use the fact that an integral must be positive if the integrand is always positive. For Responsiveness, we pass the partial derivative inside the integral in Eq. (2). For Monotonicity, we express $\hat{Q}_j(\pi) - \hat{Q}_{j'}(\pi)$ using Eq. (2) and then pair up terms to form a single integral. ■

Theorem 4 tells us that in our framework, the representative agent’s quantal response function always satisfies Monotonicity. While the stronger condition of iid payoff disturbances is not inherited by the representative agent, the weaker condition of Monotonicity is. It is this Monotonicity property that carries empirical content. In principle, subjects in an experiment could violate Monotonicity and choose actions with lower payoffs more often than actions with higher payoffs. This would be inconsistent with the predicted behavior of the representative agent.

Given our results that a representative agent’s specification of payoff noise may differ from the noise distributions in the underlying population, but that the representative quantal response function inherits regularity properties from the population, we would second Goeree, et. al.’s suggestion [4] of dealing only with regular QRE and not necessarily requiring them to be derived from a model of noisy utility maximization. Taking the regularity axioms as primitives of a stochastic choice model provides a simple and intuitive foundation, carrying empirical content and consistent

with heterogeneity.

4 Asymmetric Games

All of these results, initially presented in the context of a single population, apply to general asymmetric games. In a normal-form game with n populations of agents, we would introduce an index $i \in [1, n]$ distinguishing each population as corresponding to a particular role in the game. The strategy sets, population mixed strategy vectors, payoff functions, distributions of payoff disturbances, and aggregate quantal response functions would be distinguished for each population / role in the game by this index i . Theorem 1 would describe the existence of a distinct representative agent for each population i . Eq. (4) and (6), also indexed by i , would characterize each representative agent's distribution of payoff disturbances. And Theorems 2, 3, and 4 would apply to the representative agent from any given population.

While we obtain representative agents for each role in an asymmetric game, we caution that there is no reason to assume the existence of a single representative agent the same for all players in the game. Such an assumption would deny heterogeneity across the different roles of the game. In the special case that all players of an asymmetric game have the same number of actions, a single representative agent for all players would generate identical aggregate quantal response functions in each population.⁸ There are plenty of QRE which are incompatible with all players having

⁸In the general case of asymmetric games in which different players may have different dimensional

the same quantal response function, as the following game illustrates.

Example

Asymmetric Matching Pennies

	Left	Right
Up	9, -1	-1, 1
Down	-1, 1	1, -1

Goeree, et. al. [4] analyze the asymmetric matching pennies game shown above and find the set of possible regular QRE is the rectangle $\frac{1}{6} < p < \frac{1}{2}, \frac{1}{2} < q < 1$, where p is the probability the column player chooses left and q is the probability the row player chooses up. However, given the restriction that representative row and column players have the same quantal response function, a QRE must satisfy the additional constraint $q < 3p$ if and only if $q > 1 - p$. (See Figure 1.) This is because $q < 3p$ means that $\pi_{2R} - \pi_{2L} < \pi_{1U} - \pi_{1D}$, i.e., the cost of an error is higher for the row player, and must lead to relatively fewer errors by the row player, $q > 1 - p$. The converse holds equivalently. (Note also that if both row and column players use identical logit strategy spaces, it's not obvious exactly what is meant by a single representative agent for all players, given the fact that a representative agent does not have a representative distribution of payoff shocks. We want to have representative behavior, which requires payoff disturbances jointly dependent across actions. The problem is how to define a population average quantal response function when the individual quantal response functions have varying dimensionality. A solution is to embed lower dimensional strategy spaces in higher dimensional spaces in a natural way, by assuming there are additional strategies with infinitely negative payoffs that are never chosen.




Figure 1: On the horizontal axis is the probability the column player chooses left and on the vertical axis is the probability the row player chooses up. The entire shaded area represents the set of possible regular QRE in the given asymmetric matching pennies game when row and column players may differ. The darkened area represents the subset of these QRE consistent with the assumption that representative row and column players have the same quantal response function.

responses, the set of possible QRE is reduced to a curve extending from the center of the strategy space to the Nash Equilibrium ($p = \frac{1}{6}, q = \frac{1}{2}$) [17].) In summary, unlike the existence of a representative agent for each population, we do not necessarily have a single representative agent for all players.

5 Discussion

We have proposed a model of heterogeneous populations playing quantal response equilibria, placing no restrictions on the particular functional forms the quantal responses may take. We have found that representative agents exist for heterogeneous populations if we allow weakly admissible payoff disturbances. A representative agent chooses strategies in the same proportions as the entire population, but does not have payoff disturbances distributed in the same proportions as the population. In games with many pure strategies, representative behavior cannot arise from any iid distribution of disturbances.

This impossibility of having a representative agent with disturbances iid across actions stems from the fact that averaging probability distributions almost never preserves independence. Thus, if we believe populations of agents are heterogeneous, but desire representative-agent models, we must be willing to consider noise terms that are jointly dependent across actions. Our findings support the use of regular quantal response functions. Regular quantal response equilibrium does generate falsifiable predictions and is consistent with the representative-agent framework.

Appendix

Lemma 1 $\hat{\phi}(t)$ is continuous at $t = 0$.

Proof Recall that $|\phi(t)| \leq 1$ and $\phi(0) = 1$ for all ϕ and thus for $\hat{\phi}$ as well. We will

show for all $h > 0$, there exists $k > 0$ such that $\|t\| < k$ implies $\operatorname{Re} \left\{ \hat{\phi}(t) \right\} > 1 - h$.

Let \mathcal{K} be a compact subset of characteristic functions ϕ such that $\int_{\mathcal{K}} \Psi[\phi] d\phi > 1 - \frac{h}{4}$. Because all the ϕ are continuous at $t = 0$, we can choose $k[\phi] > 0$ such that $\operatorname{Re} \{ \phi(t) \} > 1 - \frac{h}{2}$ for all $\|t\| < k[\phi]$ and $\phi \mapsto k[\phi]$ is continuous. Then take $k = \min_{\phi \in \mathcal{K}} k[\phi]$, and $k > 0$ because the minimum is taken over a compact space and the extreme value theorem applies. We then obtain for all $\|t\| < k$,

$$\begin{aligned} \operatorname{Re} \left\{ \hat{\phi}(t) \right\} &= \int_{\phi \in \mathcal{K}} \operatorname{Re} \{ \phi(t) \} \Psi[\phi] d\phi + \int_{\phi \notin \mathcal{K}} \operatorname{Re} \{ \phi(t) \} \Psi[\phi] d\phi \\ &> \left(1 - \frac{h}{2} \right) \left(1 - \frac{h}{4} \right) + (-1) \left(\frac{h}{4} \right) \\ &= 1 - h + \frac{h^2}{8} > 1 - h. \quad \blacksquare \end{aligned}$$

Lemma 2 *The Cauchy principal value of the mean of $\hat{P}(\cdot)$ is 0. Additionally, if the random vector $(\bar{\delta}_{1j}, \dots, \widehat{\bar{\delta}}_{jj}, \dots, \bar{\delta}_{Jj})$ is distributed according to $\hat{P}(\cdot)$, then*

$$\lim_{\gamma \rightarrow \infty} \gamma \operatorname{Prob} \left\{ |\bar{\delta}_{j'j}| \geq \gamma \right\} = 0 \text{ for all } j'.$$

Proof A property of characteristic functions is that $\left[\frac{\partial}{\partial t_{j'}} \phi(t) \right]_{t=0}$ exists if and only

if:

- (i) PV $\langle \delta_{j'j} \rangle$ exists and
- (ii) $\lim_{\gamma \rightarrow \infty} \gamma \operatorname{Prob} \{ |\delta_{j'j}| \geq \gamma \} = 0$,

and when these conditions are satisfied, $\left[\frac{\partial}{\partial t_{j'}} \phi(t) \right]_{t=0} = i \operatorname{PV} \langle \delta_{j'j} \rangle$ [25, 19]. So, it suffices to show $\left[\frac{\partial}{\partial t_{j'}} \hat{\phi}(t) \right]_{t=0} = 0$ for all j' . Differentiability of $\hat{\phi}(t)$ follows from the

differentiability of all ϕ in the support of Ψ , using an argument completely analogous to the proof of continuity of $\hat{\phi}(t)$, Lemma 1. Thus, $\left[\frac{\partial}{\partial t_{j'}}\hat{\phi}(t)\right]_{t=0} = \int \Psi[\phi] \left[\frac{\partial \phi}{\partial t_{j'}}\right]_{t=0} d\phi$. For all ϕ in the support of Ψ , all j' , $\left[\frac{\partial \phi}{\partial t_{j'}}\right]_{t=0} = 0$ because $\text{PV} \langle \delta_{j'j} \rangle = 0$ and $\lim_{\gamma \rightarrow \infty} \gamma \text{Prob} \{|\delta_{j'j}| \geq \gamma\} = 0$. Each $\delta_{j'j}$ must satisfy these two conditions because the underlying ϵ_j and $\epsilon_{j'}$ are required to obey them by assumption (a3) or (a3'). ■

References

- [1] Anderson, S., De Palma, A., and Thisse, JF. (1992) *Discrete Choice Theory of Product Differentiation*. MIT Press, Cambridge, MA.
- [2] Anderson, S., Goeree, J., and Holt, C. (2004) Noisy Directional Learning and the Logit Equilibrium. *Scand. J. Econ.* 106(3), 581-602.
- [3] Fudenberg, D., and Levine, D. (1998) *The Theory of Learning in Games*. MIT Press, Cambridge, MA.
- [4] Goeree, J., Holt, C., and Palfrey, T. (2005) Regular Quantal Response Equilibrium. *Exper. Econ.* 8, 347-367.
- [5] Golman, R. (2011) Logit Equilibria with Heterogeneous Agents. Working Paper.
- [6] Golman, R. (2011) Why Learning Doesn't Add Up: Equilibrium Selection with a Composition of Learning Rules. *Int. J. Game Theory*, forthcoming.

- [7] Haile, P., Hortacsu, A., and Kosenok, G. (2008) On the Empirical Content of Quantal Response Equilibrium. *Amer. Econ. Rev.* 98(1), 180-200.
- [8] Ho, T., Wang, X., and Camerer, C. (2008) Individual Differences in EWA Learning with Partial Payoff Information. *Econ. J.* 118, 37-59.
- [9] Hofbauer, J., and Hopkins, E. (2005) Learning in Perturbed Asymmetric Games. *Games Econ. Behav.* 52(1), 133-152.
- [10] Hofbauer, J., and Sandholm, W.H. (2002) On the Global Convergence of Stochastic Fictitious Play. *Econometrica* 70, 2265-2294.
- [11] Hofbauer, J., and Sandholm, W.H. (2007) Evolution in Games with Randomly Disturbed Payoffs. *J. Econ. Theory* 132, 47-69.
- [12] Hommes, C. (2006) Heterogeneous Agent Models in Economics and Finance. In: Judd, K., Tesfatsion, L. (eds) *Handbook of Computational Economics*, vol 2, Agent-Based Computational Economics. North Holland, Amsterdam.
- [13] Hopkins, E. (1999) A Note on Best Response Dynamics. *Games Econ. Behav.* 29, 138-150.
- [14] Hopkins, E. (2002) Two Competing Models of How People Learn in Games. *Econometrica* 70, 2141-2166.
- [15] Kirman, A. (2006) Heterogeneity in Economics. *J. Econ. Interact. Coord.* 1, 89-117.

- [16] McFadden, D. (1974) Conditional Logit Analysis of Qualitative Choice Behavior. *Frontiers of Econometrics*. Academic Press, New York, NY.
- [17] McKelvey, R., and Palfrey, T. (1995) Quantal Response Equilibria for Normal Form Games. *Games. Econ. Behav.* 10(1), 6-38.
- [18] McKelvey, R., Palfrey, T., and Weber, R. (2000) The effects of payoff magnitude and heterogeneity on behavior in 2x2 games with unique mixed strategy equilibria. *J. Econ. Behav. Organ.* 42, 523-548.
- [19] Pitman, E. J. G. (1956) On the Derivatives of a Characteristic Function at the Origin. *The Annals of Mathematical Statistics* 27(4), 1156-1160.
- [20] Rogers, B., Palfrey, T., and Camerer, C. (2009) Heterogeneous Quantal Response Equilibrium and Cognitive Hierarchies. *J. Econ. Theory* 144 (4), 1440-1467.
- [21] Tsakas, E. (2005) Mixed Quantal Response Equilibria for Normal Form Games. Paper presented at the International Conference on Game Theory.
- [22] Weibull, J. (1994) The Mass-Action interpretation of Nash equilibrium. Paper presented at Nobel symposium in honor of John Nash.
- [23] Weizsäcker, G. (2003) Ignoring the Rationality of Others: Evidence from Experimental Normal-Form Games. *Games. Econ. Behav.* 44, 145-171.
- [24] Wilcox, N. (2006) Theories of Learning in Games and Heterogeneity Bias. *Econometrica* 74, 1271-1292.

- [25] Wolfe, S. (1973) On the Local Behavior of Characteristic Functions. *The Annals of Probability* 1(5), 862-866.