Quantal response equilibrium as a structural model for estimation: The missing manual*

James R. Bland
University of Toledo

Theodore L. Turocy
University of East Anglia
and Alan Turing Institute

April 21, 2023

Abstract

One of the original objectives of the (logit) quantal response equilibrium (LQRE) model was to provide a method for structural estimation of behaviour in games, when behaviour deviated from Nash equilibrium predictions. To date, only Chapter 6 of the book on quantal response equilibrium by Goeree et al. (2016) focuses on how such estimation can be implemented. We build on that chapter to provide here a more detailed treatment of the methodological issues of implementing maximum likelihood estimation of QRE. We compare the equilibrium correspondence and empirical payoff approaches to estimation, and identify some considerations in interpreting the results of those approaches when applied to the same data on the same game. We also provide a more detailed "field guide" to using numerical continuation methods to accomplish estimation, including guidance on how to tailor implementations to games with different structures.

JEL Classifications: C63, C72, C90.

Keywords: quantal response, estimation, computation, experiments.

1 Introduction

In their book surveying the history, theory, and applications of quantal response equilibrium (QRE), Goeree et al. (2016) dedicate Chapter 6 to the topic of using QRE, and in particular the logit spec-

*Please do not yet cite or re-circulate without the permission of the authors. We thank participants at the 25+ Years of QRE Conference at Caltech in March 2022, and the Behavioural Game Theory workshop at University of East Anglia, for many helpful comments and discussions. Any errors are the responsibility of the authors.

ification of QRE (LQRE), as a structural model for estimation. That chapter, which is one of the shorter chapters of the book, is given over to two principal topics. First, the chapter illustrates the conceptual process of estimating LQRE through the use of fully-explicated MATLAB code, which computes likelihood-maximising points on the LQRE correspondence for 2×2 games. Because the fitted points are guaranteed to be LQRE of the game, the chapter calls this the *equilibrium correspondence approach*. The chapter then proceeds to outline, at a very high level, some potential technical issues in computing LQRE. In light of these potential issues, the chapter introduces an alternative approach, which removes the need to solve the fixed-point problem in the definition of LQRE by substituting the empirical payoffs from the observed data; this is referred to as the *empirical payoff approach*.

In most of the chapters of Goeree et al. (2016), the authors were able to rely on citing published articles developing the theory of QRE and its application to various domains. Chapter 6 differs because there are few papers which discuss the properties of methods for using QRE as a structural model, or of computational methods for doing so, as topics in their own right. As suggested by our subtitle, "The missing manual," in this paper our goal is to fill in more completely and formally many of the details of the ideas introduced that chapter.

One of the attractive features of the development of QRE as a structural model is that it encodes information about the strategic structure of the game in the model. QRE has strong foundations in the tradition of game-theoretic analysis, as seen in the underpinning results in McKelvey and Palfrey (1995), as well as its relationship to other concepts from standard game theory, such as the purification argument of Harsanyi (1973). The encoding of the information about the strategic structure is implicated in the computational cost of computing LQRE: the limit of the set of LQRE as the precision of responses becomes infinite is a subset of the Nash equilibria of the game. This is indeed a computationally difficult problem, as we will make more precise in Section 5. The empirical payoff approach, however, in essence throws away all of this strategic information by approximating the game as a decision problem. In this paper we illustrate that the implications of this approximation are not necessarily innocuous. The empirical payoff approach produces systematically different results than the fixed point method applied to the same game and data. Differences occur across all games; in some games, most notably those with equilibria involving randomization, the results from the empirical payoff approach can be so starkly different from those obtained via the equilibrium correspondence approach as to be misleading about the fit of the data to the fully game-theoretic LQRE model.

Goeree et al. (2016) assert (p. 154) that "the computation of the logit QRE correspondence (or a selection from it) can be challenging computationally" and "in many games ... good algorithms for computing QREs are unavailable." These statements are not wrong, but neither are they precise. In particular, there is a lot of nuance hidden under the word "good." There are two different senses in which a computational problem can be "challenging" or, relatedly, an algorithm can be "good".

The method illustrated in the MATLAB codes provided in Goeree et al. (2016) indeed does not scale well. Those codes are useful as an expositional device for didactic purposes, but are not a proper foundation for production-quality code. The elegant structure of the set of LQRE as the union of one-dimensional differentiable curves makes it suitable for computation using numerical continuation methods. These methods are good not only at traversing a differentiable curve efficiently, but can find extreme points of functions, such as a likelihood function, along the curve with negligible additional computational effort.

However, in view of the computational hardness results about computing Nash equilibria, there must be some other consideration limiting the performance of numerical continuation methods, at least for some classes of games. As we develop in more detail, the key piece of information that a numerical continuation method requires is to know the direction tangent to the curve at a given point. Computing this requires knowledge of the Jacobian of the system of equations defining an LQRE. It is here where we pay the computational price for one of the attractive features of LQRE. For any finite value of the precision parameter, an LQRE has full support. This means we avoid the zero-likelihood problem in taking any empirical data to the model, as all strategies or actions occur with positive probability. However, the Jacobian of the LQRE system requires computing the derivative of the expected payoff of every strategy with respect to every other strategy, and those expected payoffs require considering all of the contingencies in the game consistent with those two strategies being played. Practical experience shows that it is the time spent in computing these payoffs which dominates the running time for LQRE computation and estimation. The structure of the Jacobian, therefore, is what allows us to formalize the assertion in Goeree et al. (2016) that computation of LQRE in some "games of incomplete information with many times and many strategies" can be slow, and therefore offer practical guidance for production-quality code to implement estimation properly using the equilibrium correspondence.

The paper is organized as follows. Section 2 recaps the essential facts about LQRE and formally states the optimization problems used by the equilibrium correspondence and empirical payoff approaches. Section 3 provides some results which partially characterize the mixed strategy profiles estimated by the empirical payoff approach, and the relationship between those profiles and the equilibrium correspondence approach. Section 4 gives a more detailed exposition of the ways in which the results of the two approaches differ by considering selected examples. Section 5 turns to the question of the theoretical and practical complexity of the equilibrium correspondence approach, outlining recommendations for improving performance to make the approach feasible in broader classes of games. Section 6 concludes with a summary discussion of the implications of the principal results.

2 The logit quantal response equilibrium

2.1 Definitions

For most of this paper, we consider n-player games in normal form. The set of players is $N = \{1, \dots, n\}$, and each player $i \in N$ has a set of J_i strategies, $S_i = \{s_{i1}, \dots, s_{iJ_i}\}$; let S be the Cartesian product of the strategy sets S_i, and $J = \sum_{i=1}^n J_i$. Let Δ_i be the set of probability distributions over S_i. An element $\pi_i \in \Delta_i$ specifies, for each strategy $s_{ij} \in S_i$, the probability π_{ij} that player i plays that strategy. A mixed-strategy profile is an element $\pi \in \Delta$, where Δ is the Cartesian product of the Δ_i. We write $\pi^0 \in \Delta$ as the **centroid** of Δ; this is the mixed-strategy profile such that $\pi^0_{ij} = \frac{1}{|S_i|}$ for all $1 \leq i \leq N$ and $1 \leq j \leq J_i$.

We define a **game** on S by specifying, for each player $i \in N$, a payoff function $u_i : S \to \mathbb{R}$. We will use suitably-decorated versions of the symbol Γ to denote such games. Two games Γ and Γ' are equal, $\Gamma = \Gamma'$, if and only if they share the same strategy space $S = S'$ and the same utility payoff function $u = u'$. In general we will work with a fixed strategy space S at any time, and so games are differentiated by their corresponding payoff function. We assume players are expected-utility maximizers; therefore the payoff of player i if a mixed-strategy profile π is played is $u_i(\pi) = \sum_{s \in S} \pi(s) u_i(s)$, where $\pi(s) = \prod_{i \in N} \pi_i(s_i)$.

Following McKelvey and Palfrey (1995), we make the following definition:

Definition 1. *Given a game, a* ***logit quantal response equilibrium (LQRE)*** *is a pair (λ, π) consisting of a real number $0 \leq \lambda < \infty$ and a mixed-strategy profile $\pi \in \Delta$ which satisfies, for all players $1 \leq i \leq n$ and all strategies $1 \leq j \leq J_i$,*

$$\pi_{ij} = \frac{\exp\left[\lambda u_{ij}(\pi)\right]}{\sum_{k=1}^{J_i} \exp\left[\lambda u_{ik}(\pi)\right]} \tag{1}$$

It will be convenient in developing some ideas, and in computation, to express the equations (1) in ratio form. It is equivalent to state that (λ, π) is an LQRE if and only if, for all players $1 \leq i \leq n$ and all strategies $1 \leq j \leq J_i - 1$,

$$\begin{array}{ll} \frac{\pi_{i,j+1}}{\pi_{ij}} = \exp\left[\lambda(u_{i,j+1}(\pi) - u_{ij}(\pi))\right] & \forall i, j : 1 \leq i \leq n, 1 \leq j \leq J_i - 1 \\ \sum_{j=1}^{J_i} \pi_{ij} = 1 & \forall i : i \leq i \leq n. \end{array} \tag{2}$$

By taking limits as the precision parameter λ becomes large, LQRE can be used to make a selection from the set of Nash equilibria.

Proposition 2 (McKelvey and Palfrey (1995), Theorem 2)**.** *Fix a game, and a sequence $\{(\lambda_i, \pi_i)\}_{t=1}^\infty$ of LQRE such that $\lim_{t\to\infty} \lambda_t = \infty$ and $\lim_{t\to\infty} \pi_t = \pi^\star$. Then $\pi^\star$ is a Nash equilibrium of the game.*

This result legitimates the following definition, as at places in our analysis it will be useful to include the limiting Nash equilibria.

Definition 3. *Given a game* Γ*, we define*

- $\mathcal{L}(\Gamma)$ *to be the (extended)* ***set of LQRE*** *of* Γ*: the union of all* (λ, π) *pairs that are LQRE, and of all pairs* $(\infty, \pi^\star)$ *where* $\pi^\star$ *is a limit of some sequence of LQRE as in Proposition 2.*

- $\mathcal{P}(\Gamma) = \{\pi : (\lambda, \pi) \in \mathcal{L}(\Gamma)\}$*, the set of mixed strategy profiles in* $\mathcal{L}(\Gamma)$*, which we refer to as the* ***locus of LQRE profiles*** *of* Γ.

Distinguishing the locus of LQRE profiles from the set of LQRE will be useful in what follows, due to the following well-known property of the logit specification.

Proposition 4. *Fix a strategy space* S *and consider two games,* Γ*, which has payoff function* u*, and* Γ'*, which has payoff function* u'*. Then* $\mathcal{P}(\Gamma) = \mathcal{P}(\Gamma')$ *if and only if* $u' = Ku + L$*, for some* $K > 0$ *and* $L \in \mathbb{R}$.

Proof. Fix a strategy profile $\pi \in \mathcal{P}(\Gamma')$. Then, there exists some λ' such that, for each $1 \leq i \leq N$ and $1 \leq j \leq J_i - 1$,

$$\frac{\pi_{i,j+1}}{\pi_{ij}} = \exp\left[\lambda'(u'_{i,j+1}(\pi) - u'_{ij}(\pi))\right] = \exp\left[K\lambda'(u_{i,j+1}(\pi) - u_{ij}(\pi))\right]. \tag{3}$$

Setting $\lambda = K\lambda'$, we have that $(\lambda, \pi) \in \mathcal{L}(\Gamma)$, and $\pi \in \mathcal{P}(\Gamma)$. □

The above calculation reminds us that the precision parameter λ is denominated in units of inverse utility. Finally, we note the following facts about the relationship between $\mathcal{L}(\Gamma)$ and $\mathcal{P}(\Gamma)$:

Fact 5. *For a given game* Γ*:*

1. *If* $\pi \in \mathcal{P}(\Gamma)$ *and* $\pi \neq \pi^0$*, then there exists a unique* λ *such that* $(\lambda, \pi) \in \mathcal{L}(\Gamma)$.

2. $(\lambda, \pi^0) \in \mathcal{L}(\Gamma)$ *for some* $0 < \lambda < \infty$ *if and only if* $(\hat{\lambda}, \pi^0) \in \mathcal{L}(\Gamma)$ *for all* $\hat{\lambda} \in \mathbb{R}_+ \cup \{\infty\}$.

2.2 Maximum-likelihood estimation

McKelvey and Palfrey (1995) introduced the convention of taking LQRE to data using maximum likelihood estimation. Given a game Γ, let $p \in \Delta$ denote observed empirical frequencies of play in a sample: p_{ij} is therefore the proportion of plays of the game in which player i played strategy

s_{ij}.[1] The best-fit LQRE is selected by maximizing the (log-)likelihood of the data over the set of LQRE,[2] where we follow the convention that $0 \log 0 = 0$.

$$\hat{\pi}(p, \Gamma) = \arg \max_{\pi \in \mathcal{P}(\Gamma)} \sum_{i=1}^{N} \sum_{j=1}^{J_i} p_{ij} \log \pi_{ij} \tag{4}$$

When this is applied directly to the full specification of a game, this is referred to in Goeree et al. (2016) as the **equilibrium correspondence approach**.

From Fact 5, the maximizer $\hat{\pi}$ maps to a unique value of λ, except in the non-generic edge case in which $\hat{\pi} = \pi^0$ and π^0 is a Nash equilibrium of the game. While we note this includes games with nontrivial strategic considerations, not least pure-coordination games, this situation is easily diagnosed. In what follows we therefore focus on the case in which either $\hat{\pi} \neq \pi^0$, or if $\hat{\pi} = \pi^0$, then π^0 is not a Nash equilibrium.

The analysis of McKelvey and Palfrey (1995) shows that, with the exception of the limiting Nash equilibrium points, $\mathcal{L}(\Gamma)$ can be expressed as the union of differentiable curves on $\mathbb{R}_+ \times \Delta$. Let $\mathcal{M}$ be the set of those curves, with typical element $m \in \mathcal{M}$. We can express the curve m in parametric form; that is, the curve is defined by some function $(\lambda(s), \pi(s))$. We can then re-express the problem (4) as

$$\hat{\pi}(p, \Gamma) = \arg \max_{m \in \mathcal{M}, (\lambda(s), \pi(s)) \in m} \sum_{i=1}^{N} \sum_{j=1}^{J_i} p_{ij} \log \pi_{ij}(s) \tag{5}$$

In this case, the first-order necessary condition for $(\lambda(s), \pi(s))$ to be a local maximizer along a given curve $m \in \mathcal{M}$ is[3]

$$\sum_{i=1}^{N} \sum_{j=1}^{J_i} p_{ij} \frac{\pi'_{ij}(s)}{\pi_{ij}(s)} = 0. \tag{6}$$

We use the symbol $\nabla \pi(s)$ to denote the vector of derivatives $\pi'_{ij}(s)$.

A few observations about our formulation (5) of the problem are in order. First, we take care to account for the existence of multiple curves comprising the feasible set. In many applications, attention is focused on the curve, referred to as the **principal branch** by McKelvey and Palfrey (1995), which contains the point $(0, \pi^0)$. Such a curve always exists in the feasible set, and fur-

[1] In applications, one would use the *counts* of plays rather than frequencies. Using counts or frequencies both result in the same likelihood-maximising LQRE, and it is these fitted LQRE which are of interest in this paper. Count data is important when the maximising log-likelihood is important, for example for likelihood-ratio tests.

[2] When McKelvey and Palfrey (1995) took LQRE to the data from Ochs (1995), which we look at as a case study below, they estimate Nash equilibrium play for one block of data. Therefore the inclusion of the limiting Nash equilibria in $\mathcal{P}(\Gamma)$ is essential for the correctness of (4).

[3] Multiple local maximizers are possible and not uncommon. For example, they occur along a one-dimensional set in the asymmetric and generalized matching pennies games in Section 4.

thermore is computationally accessible given that $(0, \pi^0)$ is known to be one endpoint of it. It may also be the most behaviorally-relevant, either because it is the unique curve that comprises the set, or because behavior is sufficiently noisy, corresponding to relatively small values of λ, that the best fit will fall on this curve. Further, we express the component curves as parameterized curves which are not (necessarily) parameterized by λ. This is essential for dealing with curves other than the principal one. McKelvey and Palfrey (1995) showed that for any given λ the number of corresponding LQRE profiles is generically odd; combined with Proposition 2 this implies that curves other than the principal one "connect" two Nash equilibria, and therefore must have at least one turning point at some finite λ.[4] It is often convenient to think of s as the arclength along the curve. However, this is not necessary; whatever scaling is used cancels out when calculating observable, empirically-relevant quantities such as the mixed-strategy profiles.

In our formulation (5) and the first-order necessary condition (6), the only place in which LQRE enters into the equation is via the set of curves $\mathcal{M}$ which comprise $\mathcal{L}(\Gamma)$. The equations therefore in principle apply to any solution concept whose values can be written as the union of a set of differentiable curves, and tell us that, to take that concept to the data, what will be important is computing the values of $\frac{\pi'_{ij}(s)}{\pi_{ij}(s)}$ along those curves. Relatedly, we note that λ does not appear explicitly in (6). The parameter λ plays an important role in determining the curves in $\mathcal{M}$; however, (6) shows we can frame the estimation problem in terms of quantities - the entries in mixed-strategy profiles - which are in principle observable. Later in the paper we will touch on the interpretation of λ values arising from maximum likelihood estimation processes.

The intuitively-simple advice that the key to solving (4) and (5) is to compute $\frac{\pi'_{ij}(s)}{\pi_{ij}(s)}$ conceals important practical computational considerations. Turocy (2005) demonstrated how a given curve $m \in \mathcal{M}$ can be followed using numerical continuation methods. An output of this procedure is the value of $\frac{\pi'_{ij}(s)}{\pi_{ij}(s)}$ along the curve, meaning that, conditional on tracing the curve, identifying local extreme points can be done at no extra computational cost.[5] However, there are both theoretical and practical reasons why the time required to solve the fixed-point problem (1) repeatedly along the curve may scale poorly in practice. We will discuss these practical considerations in greater detail in Section 5.

2.3 The empirical payoff approach

As an alternative to having to solve the LQRE fixed-point problem, Bajari and Hortacsu (2005) made the observation that, if in fact players are playing mixed strategies according to some LQRE,

[4]It is also possible to find examples where the principal curve has turning points in λ.

[5]McKelvey and Palfrey (1995) computed their estimates using a grid-search method. However, grid searches run into numerical problems for large λ. For sufficiently small λ, the LQRE correspondence is the fixed-point of a contraction mapping and therefore iterative methods could be used. However, in general iterative methods are not guaranteed to converge.

then the empirically-observed frequencies of play are consistent estimators of the probabilities assigned by the LQRE mixed-strategy profile, and, further, because in an LQRE players have correct beliefs about the play of others, are consistent estimators of players' beliefs. Therefore, the empirically-observed frequencies can be used to estimate the expected payoffs to each strategy. Given the practical relevance of the cost of computing expected payoffs when using the fixed-point approach, being able to substitute empirical expected payoffs has an obvious computational attraction.

Formally, this **empirical payoff approach** defines an **auxiliary game** $\tilde{\Gamma}(\Gamma, p)$ which is a function both of the **base game** Γ and the observed strategy frequencies $p \in \Delta$. This auxiliary game has the same set of players and strategies as the base game. In the auxiliary game, the payoff function is defined by $\tilde{u}_{ij}(\pi) = u_{ij}(p)$. This means that $\tilde{u}_{ij}$ is a *constant* function with respect to the mixed strategy profile π. As a result, an auxiliary game $\tilde{\Gamma}(\Gamma, p)$ has a trivial strategic structure; in fact it is no more than a collection of n decision problems, one for each player. Turocy (2005) pointed out that for a decision problem, the locus of LQRE profiles is exactly the path of the replicator dynamics (Taylor and Jonker, 1978) applied to the decision problem, where λ is time and the starting point is the centroid π^0.

Given a game Γ, each realized profile of empirical strategy frequencies p results in a different auxiliary game $\tilde{\Gamma}(\Gamma, p)$ and therefore a different set of LQRE in the resulting auxiliary game. Further, the set of LQRE in the realized auxiliary game differs from those in the base game. We further note that the auxiliary game in general does not preserve properties of the base game: for example, if the base game is constant-sum, the auxiliary game generally is not.

Given an auxiliary game $\tilde{\Gamma}(\Gamma, p)$, maximum likelihood estimation is then carried out as in (4). For notational compactness, for a given profile of strategy frequencies p, we will generally suppress explicit reference to the base or auxiliary games, and refer to the mixed strategy profile arising from the equilibrium correspondence approach as $\hat{\pi}^{EC}(p)$, with corresponding precision parameter $\hat{\lambda}^{EC}(p)$; those arising from the empirical payoff approach are written $\hat{\pi}^{EP}(p)$ and $\hat{\lambda}^{EP}(p)$, respectively.

3 Properties of the empirical payoff approach

The empirical payoff approach completely discards the strategic interaction of the game and treats the situation as a decision problem. This greatly simplifies the structure of the corresponding feasible set $\mathcal{M}$. As a result, we are able to state some general properties about the resulting estimator.

First, we remark that any LQRE profile in $\tilde{\Gamma}(\Gamma, p)$ can also be justified using the principle of maximum entropy.

Proposition 6. *Fix a point $\pi \in \mathcal{L}(\tilde{\Gamma}(\Gamma, p)$ and let $\overline{u} = \sum_{i=1}^{n} \sum_{j=1}^{J_i} \pi_{ij} u_{ij}(p)$ be the sum of expected payoffs across all players at profile π. Then, π maximizes entropy among all strategy profiles which yield a total payoff of at least $\overline{u}$.*

Proof. The strategy profile which maximizes entropy subject to the total payoff constraint satisfies the problem

$$\begin{aligned} \underset{\rho \in \Delta}{\text{maximize}} \quad & \sum_{i=1}^{n} \sum_{j=1}^{J_i} \rho_{ij} \log \rho_{ij} \\ \text{subject to} \quad & \sum_{i=1}^{n} \sum_{j=1}^{J_i} \rho_{ij} u_{ij}(p) \geq \overline{u} \\ & \sum_{j=1}^{J_i} \rho_{ij} = 1 \ \forall i \in N, \end{aligned}$$

where we use the fact that randomization is done independently by players, and the entropy of a product of distributions is the sum of the entropy of the individual distributions. Assigning $\lambda \geq 0$ as the Lagrange multiplier on payoffs and $\mu_i \in \mathbb{R}$ on each player i's sum-to-one constraint, we have for each player i and strategy j,

$$L_{ij} = 1 + \log \rho_{ij} - \lambda u_{ij}(p) + \mu_i = 0. \tag{7}$$

Fix player i, and two strategies s_{ij} and s_{ik} of player i. Applying (7) to both strategies and combining, we obtain

$$\log \rho_{ij} - \log \rho_{ik} = \lambda \left[u_{ij}(p) - u_{ik}(p) \right].$$

Because this is true for all players and all pairs of strategies for each player, this is equivalent to the logit probability distribution. □

Proposition 6 illustrates the interpretation of λ as the shadow price which trades off randomness and higher realised payoffs.[6] In the statement of the result, the total payoff over all players arises because of the assumption of a common λ across all players. If, as is frequently done, we use the data from all players to estimate a common value of λ, the total empirical payoff across all players determines the resulting estimate.

Proposition 7. *Fix a base game Γ and an auxiliary game $\tilde{\Gamma}(\Gamma, p)$, and suppose the maximum-*

[6]In fact, another way to derive logit is to assume that players receive utility from the act of randomizing itself, proportionally to the entropy of the mixed strategy they choose.

likelihood estimator $\hat{\pi}^{EP}(p) \neq \pi^0$. Then $\hat{\pi}^{EP}$ satisfies a conservation of total payoffs property:

$$\sum_{i=1}^{N}\sum_{j=1}^{J_i} \pi^{\star}_{ij} u_{ij}(p) = \sum_{i=1}^{N}\sum_{j=1}^{J_i} p_{ij} u_{ij}(p). \tag{8}$$

Furthermore, $\hat{\pi}^{EP}$ is unique.

Proof. To prove the first claim, $\mathcal{L}(\tilde{\Gamma}(\Gamma, p))$ consists of a single path in which all probabilities are monotonic in λ, and is the same as the path of the replicator dynamics in which λ is time. Therefore, at any $(\lambda, \pi) \in \mathcal{L}(\tilde{\Gamma}(\Gamma, p))$ for all $1 \leq i \leq N$,

$$\frac{\pi'_{ij}(\lambda)}{\pi_{ij}(\lambda)} = u_{ij}(p) - \sum_{k=1}^{J_i} \pi_{ik}(\lambda) u_{ik}(p). \tag{9}$$

Specializing this to $\pi = \hat{\pi}^{EP}(p)$ and substituting into the first-order condition (6) we obtain

$$\sum_{i=1}^{N}\sum_{j=1}^{J_i} p_{ij} \left\{ u_{ij}(p) - \sum_{k=1}^{J_i} \hat{\pi}^{EP}_{ik}(p) u_{ik}(p) \right\} = 0 \tag{10}$$

Therefore,

$$\begin{aligned} \sum_{i=1}^{N}\sum_{j=1}^{J_i} p_{ij} \sum_{k=1}^{J_i} \hat{\pi}^{EP}_{ik}(p)(\lambda) u_{ik}(p) &= \sum_{i=1}^{N}\sum_{j=1}^{J_i} p_{ij} u_{ij}(p) \\ \sum_{i=1}^{N}\sum_{j=1}^{J_i} \hat{\pi}^{EP}_{ik}(p) u_{ij}(p) &= \sum_{i=1}^{N}\sum_{j=1}^{J_i} p_{ij} u_{ij}(p), \end{aligned}$$

where the last step follows because $\sum_{j=1}^{J_i} = 1$ and $\sum_{k=1}^{J_i} \hat{\pi}^{EP}_{ik}(p) u_{ik}(p)$ is constant with respect to p_{ij}.

To establish uniqueness, the set $\mathcal{M}$ associated with the auxiliary game consists of a single curve. That curve is equivalent to the path taken by the replicator dynamics, with λ interpreted as time. Because strategy payoffs are constant, this implies the total payoff summed across all players is strictly increasing, and so any total payoff is associated with at most one point on the curve. □

Proposition 7 also identifies the conditions for which the empirical payoff approach will return an estimated profile that is not π^0: the total payoff across all players must exceed what the players would obtain by uniform randomization. Stochastic best response, in which all players play their strategies with probabilities in the same order as their rank in expected payoffs, is a sufficient condition for this to occur. It is not a necessary condition, even in 2×2 games, as subsequent examples will demonstrate.

A direct consequence of combining Proposition 6 and Proposition 7 is that we can characterize the strategy profile estimated by the empirical payoff approach.

Corollary 8. *Fix a base game* Γ *and a profile of strategy frequencies* p*. Suppose the maximum-likelihood estimator* $\hat{\pi}^{EP}_{ik}(p) \neq \pi^0$*, and let* $\overline{u} = \sum_{i=1}^{N} \sum_{j=1}^{J_i} p_{ij} u_{ij}(p)$ *be the sum of expected payoffs across players at* p*. Then* $\hat{\pi}^{EP}_{ik}(p)$ *maximizes entropy across all strategy profiles at which players attain a total payoff of at least* $\overline{u}$*.*

A justification for the empirical payoff approach is that, if the true data generating process is some strategy profile in $\mathcal{L}(\Gamma)$, then as the sample size gets large the empirical frequencies and empirical strategy payoffs will converge to the true values. If empirical strategy frequencies form an LQRE profile, then both approaches do indeed produce the same estimator.

Proposition 9. *Fix a base game* Γ *and some* $(\lambda, p) \in \mathcal{L}(\Gamma)$*. Then,* $\hat{\pi}^{EC}(p) = \hat{\pi}^{EP}(p) = p$ *and* $\hat{\lambda}^{EC}(p) = \hat{\lambda}^{EP}(p) = \lambda$*.*

Proof. As is well-known, the global, unconstrained maximizer $\hat{\pi}$ of the log-likelihood function sets $\hat{\pi}_{ij} = p$ for all $1 \leq i \leq N$ and all $1 \leq j \leq J_i$. By assumption, $p \in \mathcal{P}(\Gamma)$ and therefore $\hat{\pi}^{EC}(p) = p$. Because $p \in \mathcal{P}(\tilde{\Gamma}(\Gamma, p))$, it follows that $\hat{\pi}^{EP}(p) = p$. The definition of $\tilde{\Gamma}(\Gamma, p)$ ensures that $u_{ij}(p) = \tilde{u}_{ij}(p)$; therefore the $\hat{\lambda}$ associated with $\hat{\pi}$ is the same in both games. □

			Player 2	
			(π_2)	$(1-\pi_2)$
			L	R
Player 1	(π_1)	U	$10, 8$	$0, 18$
	$(1-\pi_1)$	D	$9, 9$	$10, 8$

Table 1: Game 1 from Selten and Chmura (2008).

To develop further the comparison between the equilibrium correspondence and empirical payoff approaches, in Table 1 we show the payoff table for Game 1 from Selten and Chmura (2008). With this and other 2×2 games below, for compactness we will write mixed-strategy profiles and empirical strategy frequencies as $\pi = (\pi_1, \pi_2)$ and $p = (p_1, p_2)$, respectively, where the first component is the probability Player 1 plays U and the second the probability Player 2 plays L. This game has a unique Nash equilibrium in mixed strategies, with $(\pi_1, \pi_2) = (.091, .909)$. Selten and Chmura report empirical frequencies of play $(p_1, p_2) = (.079, .690)$. Given this data, Figure 1 computes the locus of LQRE profiles for both the base game and the auxiliary game given p, and plots the locations of the mixed strategy profiles estimated by both approaches. The estimated precision parameters also differ, with $\hat{\lambda}^{EC} = 1.229 > 1.074 = \hat{\lambda}^{EP}$.

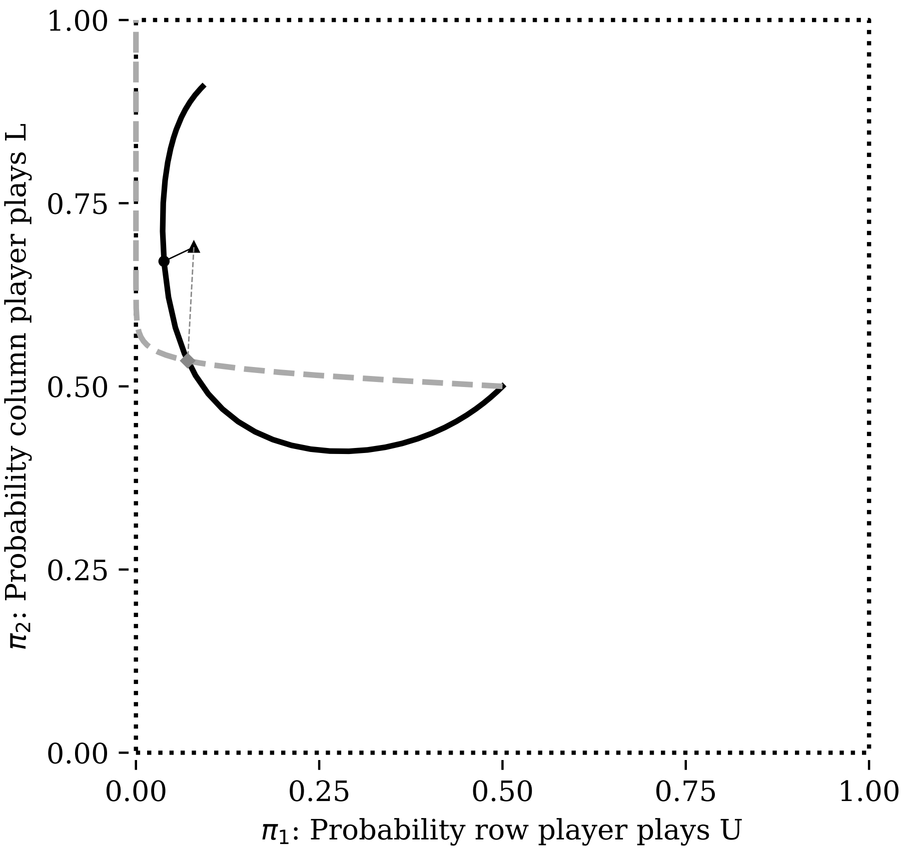

Figure 1: Data and estimation results for Game 1 from Selten and Chmura (2008). The empirical strategy probabilities are represented by the triangle. The locus of LQRE profiles of the base game is represented by the darker solid curve, and that of the auxiliary game by the lighter dashed curve. The dots indicate the corresponding maximum likelihood estimates for the two approaches.

In this initial example, we can observe that the locus of LQRE, the feasible sets over which the two approaches optimize, differ qualitatively. They do intersect at a point that is not the centroid, but the tangents to the curves at the point of intersection are quite different. It is therefore not surprising that the estimated $\hat{\pi}^{EC}$ and $\hat{\pi}^{EP}$ differ; in this case, the estimated $\hat{\pi}^{EC}$ is closer to the empirical frequencies p. Finally, in this game, it happens that $\hat{\pi}^{EP}$ is in fact an LQRE profile of the base game. In the next series of results, we will explore the extent to which these observations are general characteristics of the two approaches.

First the example shows that it is possible for the locus of LQRE of the base and auxiliary game to intersect at a point that is not the observed frequencies. However, the LQRE correspondences, inclusive of the corresponding λ values, in general do not.

Proposition 10. *Fix a base game Γ and observed strategy frequencies $p \in \Delta$. Suppose there exists a $\pi \in \Delta$, with $\pi \neq \pi^0$, such that $\pi \in \mathcal{P}(\Gamma)$ and $\pi \in \mathcal{P}(\tilde{\Gamma}(\Gamma, p))$. Let $(\lambda^{EC}, \pi) \in \mathcal{L}(\Gamma)$ and $(\lambda^{EP}, \pi) \in \mathcal{L}(\Gamma)$. Then $\lambda^{EC} = \lambda^{EP}$ if and only if $u(\pi) = u(p)$.*

Proof. Fix a player $i \in N$ and a pair of strategies $s_{ij}, s_{ik} \in S_i$. Because $(\lambda^{EC}, \pi) \in \mathcal{L}(\Gamma)$,

$$\frac{\pi_{ij}}{\pi_{ik}} = \exp\left[\lambda^{EC}(u_{ij}(\pi) - u_{ik}(\pi))\right].$$

Because $(\lambda^{EP}, \pi) \in \mathcal{L}(\tilde{\Gamma}(\Gamma, p))$,

$$\frac{\pi_{ij}}{\pi_{ik}} = \exp\left[\lambda^{EP}(u_{ij}(p) - u_{ik}(p))\right].$$

Combining these we obtain

$$\frac{\lambda^{EC}}{\lambda^{EP}} = \frac{u_{ij}(p) - u_{ik}(p)}{u_{ij}(\pi) - u_{ik}(\pi)}.$$

Therefore, $\frac{\lambda^{EC}}{\lambda^{EP}} = 1$ if and only if $u(\pi) = u(p)$. □

Proposition 10 illustrates the usefulness of thinking in terms of estimated behavior (strategy profiles) separately from the precision parameter. In general, if the two approaches agree that a mixed strategy profile is an LQRE, they will disagree on the precision parameter λ required to make that profile an LQRE. Intuitively, this occurs because λ must be interpreted in relation to (the inverse of) the units of payoff. The two approaches imply different methods for determining the payoff scale. Roughly speaking, the payoff scale used by the equilibrium correspondence approach is determined by the range of all payoffs in the game, while the empirical payoff approach uses only the realized expected payoffs of strategies to determine the size of the stakes of the game.

The domain of applicability of Proposition 10 is admittedly small. It assumes an intersection between the locus of LQRE of the base and auxiliary games exists. However, both of those are

			Player 2	
			(π_2)	$(1-\pi_2)$
			L	R
Player 1	(π_1)	U	$a_L + c_{UL}, b_U + d_{UL}$	a_R, b_U
	$(1-\pi_1)$	D	a_L, b_D	$a_R + c_{DR}, b_D + d_{DR}$

Table 2: A general 2×2 normal form game.

one-dimensional sets, and so for games other than 2×2 games, it is a non-generic situation for the two sets to intersect except at the centroid or at a limiting pure-strategy equilibrium. We therefore turn to the case of 2×2 games, where we can both develop some general results, and use those to explore the differences between the two approaches in a series of examples. Table 2 presents a generic payoff matrix for 2×2 games. Except where otherwise noted, we will focus on the case where $c_{UL} + c_{DR} \neq 0$ and $d_{UL} + d_{DR} \neq 0$. If both of these sums are zero, then in fact the base game is just a decision problem, in which the strategic incentives are independent of the other player's strategy; in that situation, the equilibrium correspondence and empirical payoff approaches are identical. Define

$$\bar{\pi}_1 = \frac{d_{DR}}{d_{UL} + d_{DR}} \text{ and } \bar{\pi}_2 = \frac{c_{DR}}{c_{UL} + c_{DR}}.$$

These are the required values of π_1 and π_2 that would make Player 2 and Player 1, respectively, indifferent between their strategies. If both $\bar{\pi}_1 \in (0, 1)$ and $\bar{\pi}_2 \in (0, 1)$, then $(\bar{\pi}_1, \bar{\pi}_2)$ is a mixed-strategy Nash equilibrium of the game. However, in the results below, $\bar{\pi}_1$ and $\bar{\pi}_2$ are permitted to correspond to improper probabilities outside the $[0, 1]$ interval.

In 2×2 games, we can specialize Proposition 10 to state precisely where an intersection between the locus of LQRE of the base and auxiliary games will occur, if it does: it will be on the line connecting (p_1, p_2) and $(\bar{\pi}_1, \bar{\pi}_2)$.

Proposition 11. *Fix a 2×2 game Γ and let (p_1, p_2) be the observed frequencies of play. If $(\pi_1, \pi_2) \in \mathcal{P}(\Gamma)$ and $(\pi_1, \pi_2) \in \mathcal{P}(\tilde{\Gamma}(\Gamma, p))$, and $(\pi_1, \pi_2) \neq p^0$ and (π_1, π_2) is not a Nash equilibrium of Γ, then*

$$\frac{\bar{\pi}_1 - p_1}{\bar{\pi}_2 - p_2} = \frac{\bar{\pi}_1 - \pi_1}{\bar{\pi}_2 - \pi_2} \tag{11}$$

Proof. Because $(\pi_1, \pi_2) \in \mathcal{P}(\Gamma)$, there exists $0 < \lambda^{EC} < \infty$ such that

$$\begin{aligned} \frac{\pi_1}{1-\pi_1} &= \exp\{\lambda^{EC}(\pi_2(c_{UL} + c_{DR}) - c_{DR})\} \\ \frac{\pi_2}{1-\pi_2} &= \exp\{\lambda^{EC}(\pi_1(d_{UL} + d_{DR}) - d_{DR})\} \end{aligned}$$

Because $(\pi_1, \pi_2) \in \mathcal{P}(\tilde{\Gamma}(\Gamma, p))$, there exists $0 < \lambda^{EC} < \infty$ such that

$$\begin{aligned}
\frac{\pi_1}{1-\pi_1} &= \exp\{\lambda^{EC}(p_2(c_{UL}+c_{DR})-c_{DR})\} \\
\frac{\pi_2}{1-\pi_2} &= \exp\{\lambda^{EC}(p_1(d_{UL}+d_{DR})-d_{DR})\}
\end{aligned}$$

Combining the above,

$$\begin{aligned}
\frac{\lambda^{EC}}{\lambda^{EP}} = \frac{\pi_2(c_{UL}+c_{DR})-c_{DR}}{p_2(c_{UL}+c_{DR})-c_{DR}} &= \frac{\pi_1(d_{UL}+d_{DR})-d_{DR}}{p_1(d_{UL}+d_{DR})-d_{DR}} \\
\frac{\pi_2 - \frac{c_{DR}}{c_{UL}+c_{DR}}}{p_2 - \frac{c_{DR}}{c_{UL}+c_{DR}}} &= \frac{\pi_1 - \frac{d_{DR}}{d_{UL}+d_{DR}}}{p_1 - \frac{d_{DR}}{d_{UL}+d_{DR}}} \\
\frac{\bar{\pi}_1 - p_1}{\bar{\pi}_2 - p_2} &= \frac{\bar{\pi}_1 - \pi_1}{\bar{\pi}_2 - \pi_2}.
\end{aligned}$$

□

Proposition 11 is illustrated by Figure 1. In that example, the observed data, the intersection of the two loci of LQRE, and the mixed-strategy equilibrium are all collinear. The Proposition confirms this is a general characteristic of 2×2 games. However, in Figure 1 the intersection of the two loci is the profile $\hat{\pi}^{EP}$ estimated by the empirical payoff method. To address how general this occurrence is, the next Proposition provides a characterisation of the estimated profile $\hat{\pi}^{EP}$.

Proposition 12. *Fix a 2×2 game Γ and observed strategy frequencies (p_1, p_2). Then, $(\hat{\pi}_1^{EP}, \hat{\pi}_2^{EP})$ satisfies*

$$\frac{p_1 - \hat{\pi}_1^{EP}}{p_2 - \hat{\pi}_2^{EP}} = -\frac{d_{UL}+d_{DR}}{c_{UL}+c_{DR}} \times \frac{p_1 - \tilde{\pi}_1}{p_2 - \tilde{\pi}_2}. \tag{12}$$

Proof. In the auxiliary game, monotonicity of the unique curve comprising the set of LQRE allows us to parameterize the curve as $(\pi_1(\lambda), \pi_2(\lambda))$. The empirical payoff estimator $(\hat{\pi}_1^{EP}, \hat{\pi}_2^{EP})$ solves

$$p_1\frac{\pi_1'(\lambda)}{\pi_1(\lambda)} - (1-p_1)\frac{\pi_1'(\lambda)}{1-\pi_1(\lambda)} + p_2\frac{\pi_2'(\lambda)}{\pi_2(\lambda)} - (1-p_2)\frac{\pi_2'(\lambda)}{1-\pi_2(\lambda)} = 0.$$

We can re-write this as

$$\frac{p_1 - \pi_1(\lambda)}{p_2 - \pi_2(\lambda)} = -\frac{\pi_2'(\lambda)}{\pi_2(\lambda)[1-\pi_2(\lambda)]} \cdot \frac{\pi_1(\lambda)[1-\pi_1(\lambda)]}{\pi_1'(\lambda)}.$$

Let $\Delta_1 = u_U(p) - u_D(p) = p_2(c_{UL}+c_{DR}) - c_{DR}$ and $\Delta_2 = u_L(p) - u_R(p) = p_1(d_{UL}+d_{DR}) - d_{DR}$.

Then, for $i = 1, 2$, $\pi_i(\lambda) = \frac{e^{\lambda\Delta_i}}{1+e^{\lambda\Delta_i}}$, and $\pi_i'(\lambda) = \Delta_i \pi_i(\lambda)(1 - \pi_i(\lambda))$. Therefore,

$$\begin{aligned} \frac{p_1 - \pi_1(\lambda)}{p_2 - \pi_2(\lambda)} &= -\frac{\Delta_2}{\Delta_1} = -\frac{p_1(d_{UL} + d_{DR}) - d_{DR}}{p_2(c_{UL} + c_{DR}) - c_{DR}} \\ &= -\frac{d_{UL} + d_{DR}}{c_{UL} + c_{DR}} \times \frac{p_1 - \bar{\pi}_1}{p_2 - \bar{\pi}_2}. \end{aligned}$$

□

Corollary 13. *Let Γ be a 2×2 game, and let (p_1, p_2) be the observed frequencies of play. Suppose $(\pi_1, \pi_2) \in \mathcal{P}(\Gamma)$ and $(\pi_1, \pi_2) \in \mathcal{P}(\tilde{\Gamma}(\Gamma, p))$, and $(\pi_1, \pi_2) \neq p^0$ and (π_1, π_2) is not a Nash equilibrium of Γ. Then $(\hat{\pi}_1^{EP}, \hat{\pi}_2^{EP}) = (\pi_1, \pi_2)$ if and only if $d_{UL} + d_{DR} = -(c_{UL} + c_{DR})$. In particular, $(\hat{\pi}_1^{EP}, \hat{\pi}_2^{EP}) = (\pi_1, \pi_2)$ for constant-sum games.*

Proof. Applying $d_{UL} + d_{DR} = -(c_{UL} + c_{DR})$ to (3) from Proposition 12, $\frac{p_1 - \pi_1}{p_2 - \pi_2} = -\frac{\bar{\pi}_1 - p_1}{\bar{\pi}_2 - p_2}$. The claim then follows by applying Proposition 11. To complete the proof, observe that $d_{UL} + d_{DR} = -(c_{UL} + c_{DR})$ is necessarily satisfied by any constant-sum game. □

Game 1 in Table 1 satisfies $d_{UL} + d_{DR} = -(c_{UL} + c_{DR})$, and therefore $\hat{\pi}^{EP}$ lies at the intersection of the two sets of LQRE profiles. Indeed, all games considered by Selten and Chmura (2008) satisfy this property. However, this provides another cautionary note about comparing values of $\hat{\lambda}^{EC}$ and $\hat{\lambda}^{EP}$. Under the conditions of Corollary 13, $\hat{\pi}^{EP}$ is in fact an LQRE profile of the base game; however, $(\hat{\lambda}^{EP}, \hat{\pi}^{EP})$ is **not** an LQRE of the base game!

4 Some illustrative examples

To build understanding of the properties of the empirical payoff approach compared to the equilibrium correspondence approach, we look at several 2×2 games. These games are simple enough to be amenable to graphical analysis to help unpack the different effects at work in distinguishing the operation of the empirical payoff approach from the equilibrium correspondence approach. At the same time, even this simple class of games is rich enough to capture the key considerations.

4.1 Asymmetric matching pennies

McKelvey and Palfrey (1995) illustrated the usefulness of LQRE using data from an asymmetric version of matching pennies studied by Ochs (1995). The payoff table of the version they considered is in Table 3b. This game has a unique mixed strategy equilibrium at $\pi^\star = (\pi_1^\star, \pi_2^\star) = (0.5, 0.2)$. The locus of LQRE profiles has a U-shape: in all LQRE, the row player plays U more than half the time. It is this feature of the LQRE correspondence in this game which illustrates

LQRE's usefulness as a model which can explain the "own-payoff" effects observed in experimental studies on this and similar games. In addition, the analysis of McKelvey and Palfrey (1995) of the data reported in Ochs (1995) results in an estimate of $\lambda = \infty$ - that is, Nash equilibrium - for the final periods of play.

			Player 2 (π_2) L	Player 2 $(1-\pi_2)$ R
Player 1	(π_1)	U	$4, 0$	$0, 1$
	$(1-\pi_1)$	D	$0, 1$	$1, 0$

(a) Payoff table in original units from Ochs (1995).

			Player 2 (π_2) L	Player 2 $(1-\pi_2)$ R
Player 1	(π_1)	U	$1.1141, 0.0000$	$0.0000, 1.1141$
	$(1-\pi_1)$	D	$0.0000, 1.1141$	$0.2785, 0.0000$

(b) Payoff table as transformed by McKelvey and Palfrey (1995).

	Actual		Fixed-point			Empirical payoff		
Periods	p_1	p_2	$\pi_1^\star$	$\pi_2^\star$	$\lambda^\star$	$\pi_1^\star$	$\pi_2^\star$	$\lambda^\star$
1-16	0.527	0.366	0.615	0.383	1.856	0.557	0.485	1.007
17-32	0.573	0.393	0.610	0.405	1.568	0.601	0.439	1.518
33-48	0.610	0.302	0.614	0.301	3.306	0.617	0.383	3.344
49-52	0.455	0.285	0.500	0.200	∞	0.500	0.500	0.000

(c) Actual data and results of LQRE estimation, with λ expressed in transformed units.

Table 3: Replication of McKelvey and Palfrey (1995) analysis of Ochs (1995), Game 3. Our estimates using the equilibrium correspondence approach match those reported by McKelvey and Palfrey. We add estimates using the empirical payoff approach.

McKelvey and Palfrey (1995) broke down the 52 periods of play in the data into four blocks, and report separate fits for each block. We replicate this analysis in Table 3c, where using the equilibrium correspondence approach we replicate exactly the same fitted LQRE mixed strategy profiles and corresponding values of λ. We additionally compute the LQRE mixed strategy profiles and corresponding values of λ using the empirical payoff approach. For the middle two blocks, periods 17-32 and 33-48, the two approaches return very similar results. For early play, periods 1-16, the approaches differ noticeably, with the empirical payoff approach reporting a smaller λ and a mixed-strategy profile closer to the centroid, representing an assessment of play that is less precise. In the final block, periods 49-52, the contrast between the two approaches is maximally stark. The equilibrium correspondence approach suggests play has converged to Nash, while the empirical payoff approach estimates that play is best described by uniform randomization.

We can build intuition for what drives the similarities, and more importantly the differences, of the two methods via a graphical analysis, which we provide in Figure 2. Each panel corresponds to one block of periods. The middle blocks, periods 16-32 and 33-48, are ones in which empirical play is close to LQRE of the base game. In such a situation, the loci of LQRE profiles of the base game and the auxiliary game are close to each other, and indeed will intersect nearby. Play in periods 1-16 is farther from an LQRE of the base game. We see the resulting estimates differ substantially between the two methods. The plots also allow us to recognise what causes the two approach to differ so greatly in the final block of periods. This game has the property that all base game LQRE lie in the quadrant of the space of mixed strategies in which the pure strategy profile (U, R) is the best response. In the first three blocks, empirical play lies in this quadrant, and so play is stochastically consistent with best response in that the higher-payoff strategy is played more frequently by both players. For any empirical frequencies of play which lie in this quadrant, the locus of LQRE profiles of the corresponding auxiliary game will connect the centroid to (U, R). In periods 49-52, the row player plays U less than one-half of the time, and therefore the strict best response for the column player is to play L. The locus of the LQRE profiles of the auxiliary game is therefore quite different; because the column player plays L (much) less than one-half of the time, the likelihood-maximizing point on that locus is the centroid. In light of Proposition 7, we observe also that this means that, despite the empirical strategy frequencies in the final block being not very far from the Nash equilibrium probabilities, the total expected payoff of the two players is actually less than they would get by both randomizing uniformly.

The panels of Figure 2 show that the analysis of the performance of the empirical payoff approach is complicated by the fact that the auxiliary game, and therefore the auxiliary game's LQRE, are functions of the observed data. To organize our explorations, we can use use Proposition 4, which showed that two games have the same locus of LQRE profiles if their payoff functions are affine transformations of each other. Given any empirical frequency of play p, the payoffs to each strategy in the corresponding auxiliary game $\tilde{\Gamma}(p)$ depend only on the p of the other player, but not the strategy the other player plays in the auxiliary game. Therefore, it follows that, among auxiliary games, the locus $\mathcal{P}(\tilde{\Gamma}(p))$ is constant for all p such that $\frac{u_U(p)-u_D(p)}{u_L(p)-u_R(p)} = K$ for some K. Because payoffs in two-player games are linear in the probabilities of the other player's mixed strategy, these sets of p are rays emanating from the mixed strategy equilibrium.

To simplify our further calculations we use the original payoff matrix as in Table 3a. It is straightforward to verify that the mixed-strategy profile $\pi = (\pi_1, \pi_2) = \left(\frac{5}{7}, \frac{2}{7}\right)$ is an LQRE, with $\lambda = \frac{7}{3}[\ln 5 - \ln 2]$.[7] At this profile, $u_U(\pi) - u_D(\pi) = \frac{3}{7}$ and $u_L(\pi) - u_R(\pi) = -\frac{3}{7}$. From this, it

[7]The analysis of Friedman and Mauersberger (2022) provides a method which makes it possible to find, in many

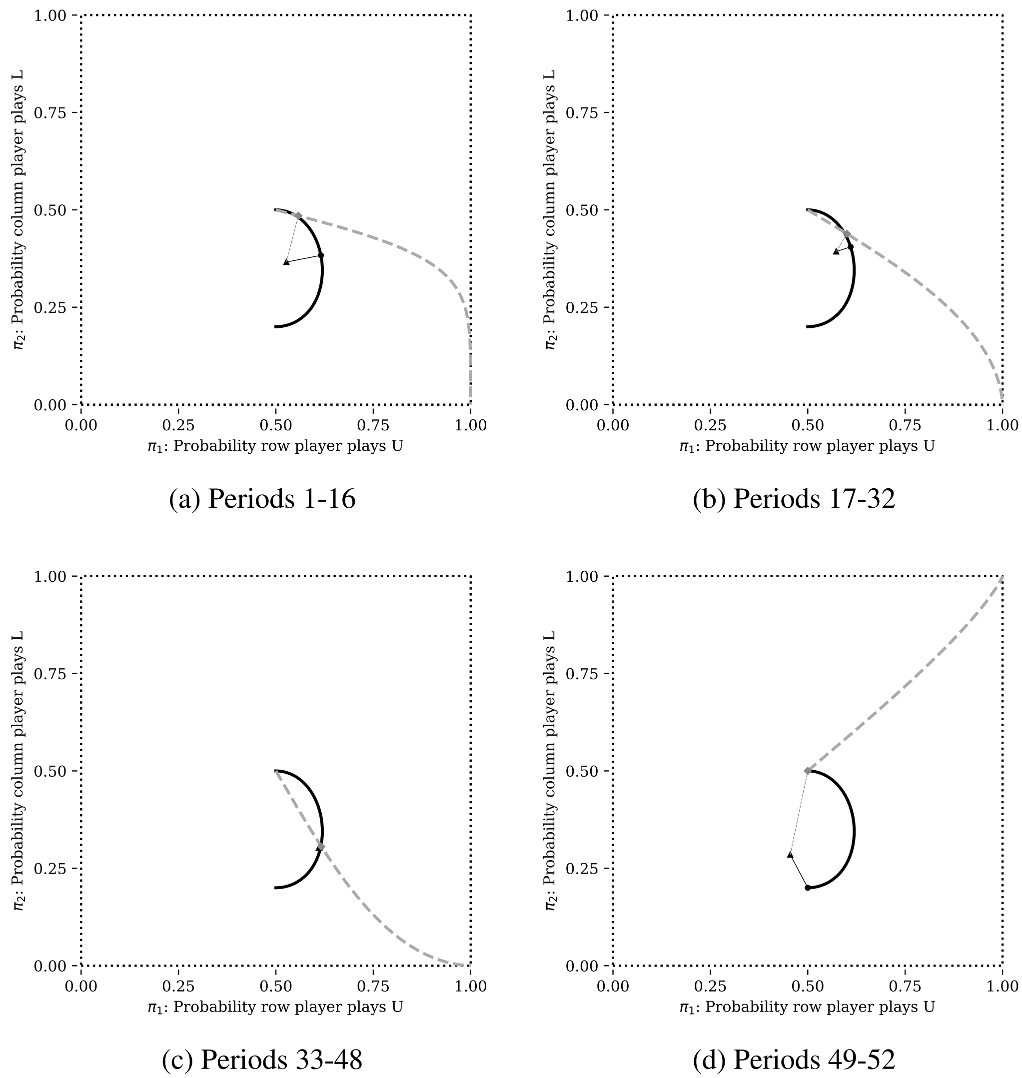

(a) Periods 1-16

(b) Periods 17-32

(c) Periods 33-48

(d) Periods 49-52

Figure 2: Graphical comparison of LQRE analysis of Table 3 via the equilibrium correspondence and empirical payoff approaches. The solid curve is the locus of LQRE profiles of the base game. The triangle indicates the location of the empirical frequencies of play. Each dashed curve represents the locus of LQRE profiles of the auxiliary game arising from the empirical data. Dotted lines connect the empirical data to the fitted profiles on the loci of the respective approaches.

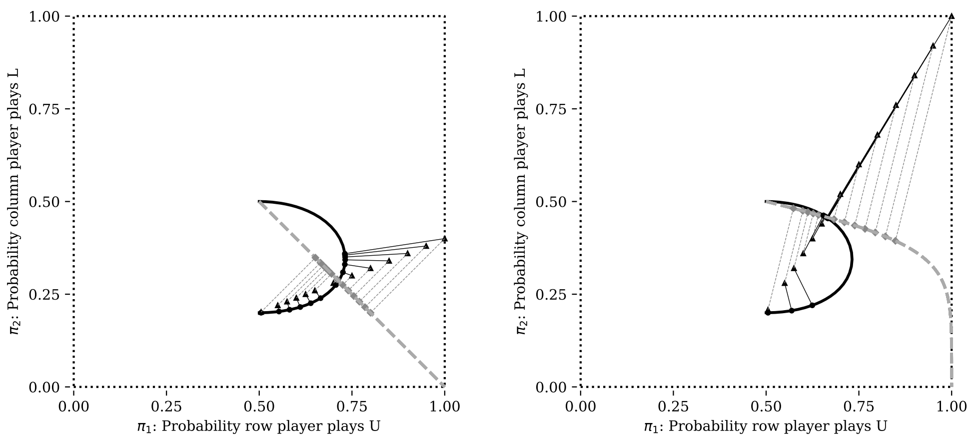

(a) Observed data $\frac{u_U(p)-u_D(p)}{u_L(p)-u_R(p)} = -1$. (b) Observed data $\frac{u_U(p)-u_D(p)}{u_L(p)-u_R(p)} = -4$.

Figure 3: Graphical comparison of LQRE analysis of Table 3 via the equilibrium correspondence and empirical payoff approaches. In each panel, we consider families of empirical data which result in auxiliary games with the same locus of LQRE profiles.

follows that the locus of LQRE for the auxiliary game $\tilde{\Gamma}(\Gamma, \pi)$ is a straight line segment between the centroid and $(1, 0)$. By Proposition 4, this is also the the locus of LQRE for all auxiliary games $\tilde{\Gamma}(\Gamma, p)$ where the observed strategy frequencies p lie on the line segment between $\left(\frac{2}{5}, 1\right)$ and the Nash equilibrium $\left(\frac{1}{5}, \frac{1}{2}\right)$. In Figure 3a we summarize how the estimates returned by the two approaches compare along this line segment. As play moves towards Nash, the equilibrium correspondence approach returns estimates which move along the locus of LQRE profiles in the direction of more precise play; in contrast, the empirical payoff approach returns estimates which move in the direction of *less* precise play. This contrast is not unique to this ray of observed strategy frequencies, but is a general feature of this game. In Figure 3b we repeat the exercise along the ray of strategy frequencies lying on the line segment between $(1, 1)$ and $\left(\frac{1}{5}, \frac{1}{2}\right)$. The same pattern emerges: as play moves towards Nash, the equilibrium correspondence approach moves along the locus towards more precise play, while the empirical payoff approach moves in the direction of less precise play.

The preceding analysis focuses on the quadrant of mixed strategies in which $p_1 \geq \frac{1}{2}$ and $p_2 \geq \frac{1}{5}$, which is the quadrant which contains the LQRE locus of the game and in which the observed strategy frequencies from the first three blocks of periods are located. The best reply to any strategy profile located in this quadrant is (U, L). Therefore, the locus of LQRE of the auxiliary game generated by any strategy profile in this quadrant is a curve which monotonically connects the centroid to $(1, 0)$, and cuts across the locus of LQRE of the base game at some point; the shape of the auxiliary game locus of LQRE and the point of intersection depend on the details

two-player games, strategy profiles with rational probabilities which are LQRE.

of the observed strategy frequencies.

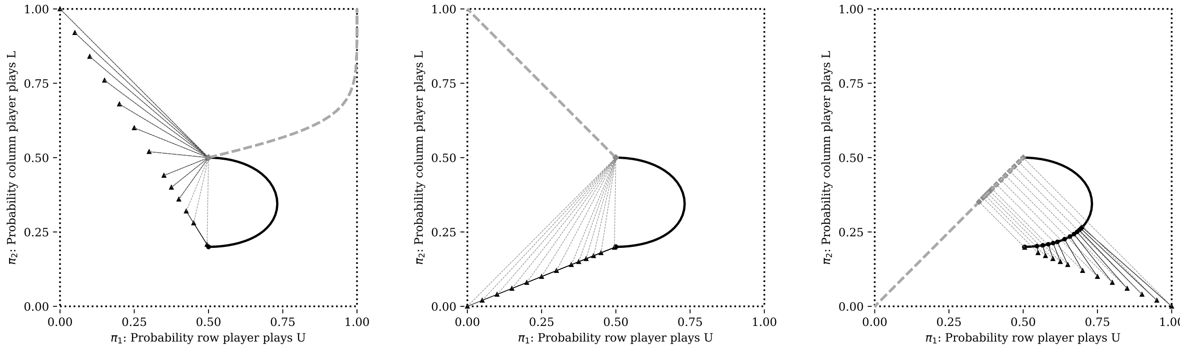

(a) Best response (D, L). (b) Best response (D, R). (c) Best response (U, R).

Figure 4: Graphical comparison of LQRE analysis of Table 3 via the equilibrium correspondence and empirical payoff approaches. Each panel corresponds to a quadrant of the set of mixed strategies in which one or both strategy frequencies is inconsistent with stochastic best response. In each panel, we consider families of empirical data which result in auxiliary games with the same locus of LQRE profiles.

The data from the fourth block of periods, which produced the maximally contrasting results of Nash play from the equilibrium correspondence approach and uniform play from the empirical payoff approach, falls outside this quadrant. In the other three quadrants, the locus of LQRE of the auxiliary game does not intersect the locus of LQRE of the base game, except at the centroid. To illustrate this, in Figure 4 we repeat the exercise from Figure 3 while looking at line segments of observed strategy frequencies falling in each of the other three quadrants. Figure 4a looks at the quadrant in which the data from the final block of periods occurs. In this quadrant, the pure-strategy best response is (D, L), and therefore the locus of LQRE of the auxiliary game, moves away from the quadrant in both dimensions as precision increases. Therefore, the empirical payoff approach maps any empirical strategy frequencies falling in this quadrant to maximally imprecise play at the centroid. This is true even for any empirical strategy frequency in this quadrant arbitrarily close to the Nash equilibrium. This illustrates the mechanics behind the behavior of the empirical payoff approach on the final block of data, and that it is a generic phenomenon that will occur with substantial probability even if play is arbitrarily close to the Nash equilibrium. Figure 4b looks at the quadrant in which (D, R) is the best response profile to the observed strategy frequencies. As illustrated by the example ray, the equilibrium correspondence approach maps all points in this quadrant to the mixed-strategy equilibrium. In contrast, the empirical payoff approach maps all points to the maximally-imprecise strategy profile at the centroid.

Finally, Figure 4c looks at the quadrant in which (U, R) is the best response profile. Along the sample ray of observed strategy frequencies shown, as the frequencies move towards the mixed-strategy equilibrium, the equilibrium correspondence approach estimates move along the locus of

			Player 2	
			(π_2)	$(1-\pi_2)$
			L	R
Player 1	(π_1)	U	$4, 0$	$0, 12$
	$(1-\pi_1)$	D	$0, 15$	$15, 0$

Table 4: A generalized matching pennies game.

LQRE profiles towards the equilibrium as well. In contrast, the estimates from the empirical payoff approach limit to the mixed strategy profile $(.35, .35)$. This mixed strategy profile is the closest point on the auxiliary game locus of LQRE to the mixed-strategy equilibrium. However, recall that the mapping between the locus of LQRE of the auxiliary game and values of the precision parameter λ are a function of the scale of the difference in the strategy values. For observed strategy frequencies close to the mixed-strategy equilibrium, the differences in strategy values $u_U - u_D$ and $u_L - u_R$ become small. Therefore, for any point on the locus of LQRE of the auxiliary game, the corresponding λ becomes larger. Intuitively, this is because any profile on the locus of LQRE that is not the centroid implies that responses are very precise, and therefore λ is large. So, along this ray of observed frequencies, the value of λ estimated by the empirical payoff approach does tend to infinity, even though the fitted mixed strategy profile converges to an interior point that is not the mixed-strategy equilibrium. As a further remark, this illustrates that although in this quadrant the observed strategy frequencies are not consistent with stochastic best response to each other, nevertheless the estimated strategy profile is not the centroid, but instead pick out a point on the locus of LRE of the auxiliary game which have the same total payoff to both players, as in Proposition 7.

4.2 Generalized matching pennies

It is fitting to open the examples with the Ochs asymmetric matching pennies game, as it was the first 2×2 game to be analyzed through the lens of LQRE. However, the LQRE correspondence of that game does have a few special features. The locus of LQRE profiles is always contained in the same quadrant relative to the mixed-strategy equilibrium, and the equilibrium involves one player mixing equiprobably over their strategies. This latter feature made the game ideal for showcasing the own-payoff effects that LQRE can capture, but also restricts the possibilities for the comparative statics of the empirical payoff method as a function of the observed strategy frequencies.

Predictions of LQRE have been explored further in the broader class of games dubbed "generalized matching pennies" by Goeree et al. (2003). Table 4 presents a parameterization of a generalized matching pennies game we have chosen to draw some contrasts to the Ochs game.

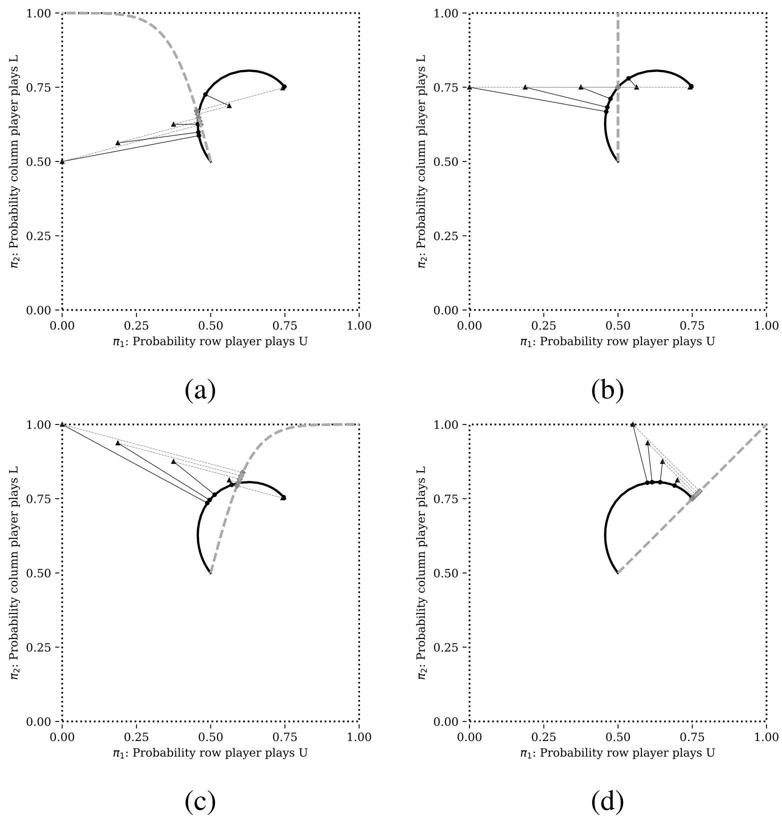

Figure 5: Four examples of comparative statics of $\hat{\pi}^{EC}$ and $\hat{\pi}^{EP}$ for the generalized matching pennies game in Table 4.

This game has its unique equilibrium at $(\pi_1, \pi_2) = \left(\frac{3}{4}, \frac{3}{4}\right)$. Figure 5a through Figure 5c display a parallel exercise to that in Figure 3. Each panel illustrates the estimates of $\hat{\pi}^{EC}$ and $\hat{\pi}^{EP}$ for a set of empirical strategy frequencies that share a common locus of LQRE profiles in the auxiliary game. Figure 5a demonstrates a case in which, as the observed strategy frequencies move towards the mixed-strategy equilibrium, both $\hat{\pi}^{EC}$ and $\hat{\pi}^{EP}$ move along their respective loci in the direction of more precise responses. In this case, the locus of LQRE profiles of the auxiliary game is not a bad approximation of the locus in the base game for a range in the vicinity of the centroid, so estimates are similar in the two approaches for the less precise points shown. The approaches diverge for more precise play, with the empirical payoff approach limiting to a point correponding to rather imprecise play. Figure 5c demonstrates a reversal of this, with the pattern observed in the Ochs game of more precise play being mapped by the empirical payoff approach to less precise estimates. Figure 5b shows where the switching point lies: in this case as play becomes more precise $\hat{\pi}^{EC}$ also becomes more precise while $\hat{\pi}^{EP}$ is constant. This shows that even in the simplest case of 2×2 games, the comparative statics of $\hat{\pi}^{EP}$ are not just game-dependent but even data-dependent within the game.

In Figure 5d we show the case in which the locus of LQRE profiles of the auxiliary game intersects that of the base game at the equilibrium point itself. This occurs for any empirical strategy probabilities on the line segment between $\left(\frac{11}{20}, 1\right)$ and the equilibrium. For these points, the estimates of both approaches do converge to the Nash equilibrium as play becomes more precise; however they do so from opposite directions, because the loci of LQRE profiles of the two approaches are orthogonal to each other at the equilibrium point. This line segment demonstrates the *only* one along which the empirical payoff approach can successfully estimate a profile that approximates a Nash equilibrium. Furthermore, $\hat{\lambda}^{EP}$ also converges to infinity along this line segment, even though the estimated point is precisely in the middle of the locus of LQRE profiles. This is a consequence of the empirical payoff approach using the empirical payoffs to determine the scale of payoffs in the game, and therefore the scale of λ, along the locus.

4.3 Estimation of λ as a function of observed data

Our analysis has focused principally on the locus of LQRE. In this game, there is a one-to-one mapping between λ and elements of the locus of LQRE, and so therefore it is meaningful to say that moving along the locus from the centroid towards the mixed-strategy equilibrium is a direction of increasing precision. However, this ordinal interpretation only applies within a given locus of LQRE; with the exception of endpoints, we cannot rank whether a point on the locus of LQRE for the base game is more or less precise than a point on the locus of LQRE for an auxiliary game. In taking LQRE to the data, it is customary to report estimates of λ itself.

We return to the Ochs asymmetric matching pennies game with payoffs as in Table 3a. To compare the λ values fitted by the two approaches, in Figure 6 we plot the level sets of $\hat{\lambda}^{EC}$ and $\hat{\lambda}^{EP}$ for selected values of λ.[8]

From this analysis, we see that for most realizations of the empirical strategy frequencies p, we will have $\hat{\lambda}^{EC} > \hat{\lambda}^{EP}$. There are two regions which are exceptions. For λ less than about 1, there is a region above and outside the locus of LQRE profiles where $\hat{\lambda}^{EP} > \hat{\lambda}^{EC}$. For λ greater than about 1, there likewise is a region "inside" the C-shape of the locus where this occurs. There are therefore two distinct factors at work which jointly determine whether $\hat{\lambda}^{EP}$ is greater than or less than $\hat{\lambda}^{EC}$. First, the level sets of $\hat{\lambda}^{EC}$ are always straight line segments.[9] In contrast, the level sets of $\hat{\lambda}^{EP}$ are curved, and specifically in the case of this game the upper-contour sets of $\hat{\lambda}^{EP}$ are convex. The curvature of the level sets of $\hat{\lambda}^{EP}$ is a consequence of the locus of LQRE profiles of the auxiliary game changing as a function of the observed data. Second, although Proposition 9

[8]The level sets for $\hat{\lambda}^{EC}$ converge along a line segment. Points along this line segment are ones in which the likelihood function has two local maximizers, one corresponding to π_2 closer to $\frac{1}{2}$, and the other corresponding to π_2 closer to the equilibrium probability of $\frac{1}{5}$. The jog in the line segment in this region is an artefact of the plotting algorithm struggling with the level sets meeting and then ending along this line segment.

[9]This is true in all games: for all games, these level sets are subsets of a hyperplane.

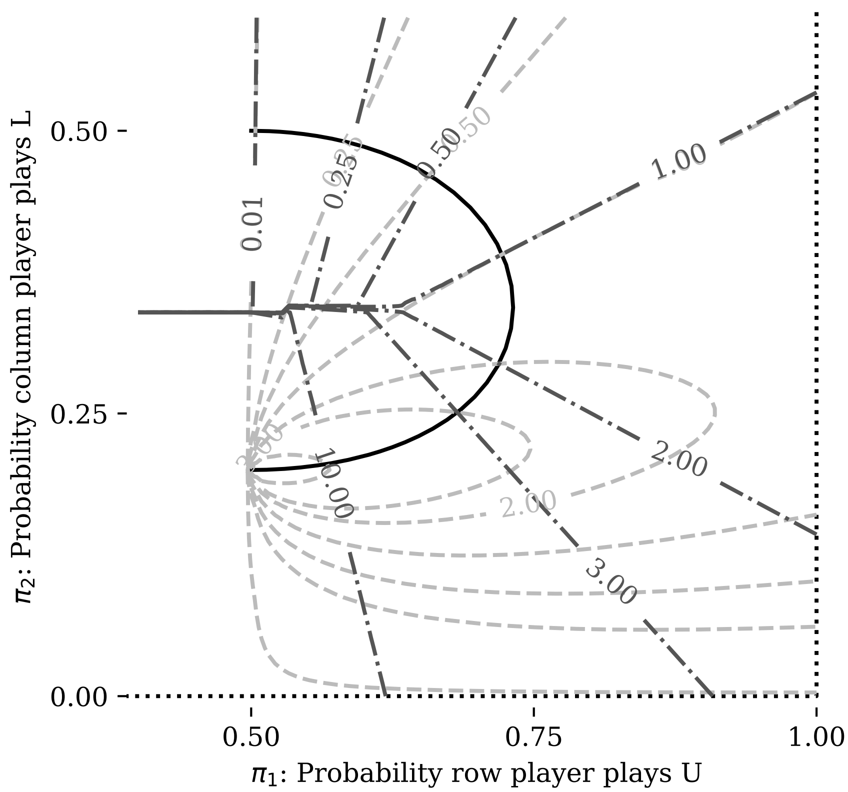

Figure 6: Level sets of $\hat{\lambda}^{EC}$ and $\hat{\lambda}^{EP}$ for various values of λ for the Ochs game in Table 3a. The solid line is the locus of LQRE profiles for the base game. The darker dot-dashed lines correspond to level sets of $\hat{\lambda}^{EC}$, and lighter dashed lines to level sets of $\hat{\lambda}^{EP}$.

ensures that the level sets with the same λ value must meet at a point on the base game locus of LQRE corresponding to that λ, in this game they are not tangent to each other but cross. The curvature of the level sets accounts for the general tendency for $\hat{\lambda}^{EC} > \hat{\lambda}^{EP}$. The regions of $\hat{\lambda}^{EP} > \hat{\lambda}^{EC}$, and the fact that these regions sit along one side or the other of the locus of LQRE profiles, is a consequence of the crossing of the level sets coming from the two approaches.

			Player 2	
			(π_2)	$(1-\pi_2)$
			L	R
Player 1	(π_1)	U	$a_L + c_{UL}, b_U + c_{UL}$	a_R, b_U
	$(1-\pi_1)$	D	a_L, b_D	$a_R + c_{DR}, b_D + c_{DR}$

Table 5: A family of 2×2 normal form games where all LQRE are symmetric. This is obtained from the payoffs in Table 2 by setting $d_{UL} = c_{UL}$ and $d_{DR} = c_{DR}$.

The crossing of the level sets for a given λ is a consequence of the crossing of the locus of LQRE of the base game and the auxiliary game, which we observed in Figure 3. Understanding the interplay of the two factors contributing to whether $\hat{\lambda}^{EC} > \hat{\lambda}^{EP}$ is challenging because, unlike in Figure 3, we cannot only look at convenient line segments along which the locus of LQRE profiles of the auxiliary game does not change. We can focus on the curvature of the level sets of $\hat{\lambda}^{EP}$ by looking at a family of games in which the level sets of $\hat{\lambda}^{EC}$ and $\hat{\lambda}^{EP}$ for a given λ are tangent when crossing the base game LQRE locus. To do this, we need a family of games in which the locus of LQRE is the same for both the base game and all auxiliary games, for realized strategy frequencies which are LQRE of the base game. We accomplish this by looking at 2×2 games where all LQRE are symmetric. Table 5 specializes the payoff structure of Table 2 to this case. In this family of games, $\Delta_1(\pi) \equiv u_U(\pi) - u_D(\pi) = \pi_2(c_{UL} + c_{DR}) - c_{DR}$ is the row player's payoff premium to playing U, and $\Delta_2(\pi) \equiv u_L(\pi) - u_R(\pi) = \pi_1(c_{UL} + c_{DR}) - c_{DR}$ is the column player's payoff premium to playing L. Without loss of generality we will assume $c_{UL} > 0$, so (U, L) is always a Nash equilibrium, and also assume $c_{UL} > c_{DR}$. When $c_{DR} > 0$, (D, R) is also a Nash equilibrium, and the locus of LQRE profiles consists of two curves: one, the principal branch, is the line segment connecting the centroid to (U, L); the other is the line segment connecting the mixed-strategy equilibrium to (D, R). When $c_{DR} \le 0$, the locus of LQRE profiles consists only of the line segment connecting the centroid to (U, L). To streamline the analysis we will focus on the quadrant containing the principal branch, that is, where $\pi_1 > \frac{1}{2}$ and $\pi_2 > \frac{1}{2}$; the key observations apply as well to the other branch, when it exists, with a suitable extension of notation.

Given observed strategy frequencies (p_1, p_2) in this quadrant, the equilibrium correspondence approach is equivalent to solving the problem $\max_x (p_1 + p_2) \log \pi + (2 - p_1 - p_2) \log(1 - \pi)$, which has solution $\hat{\pi}^{EC} = \frac{p_1 + p_2}{2}$. It is therefore natural to express any observed frequencies (p_1, p_2) as

$(p + \delta, p - \delta)$, where p is an LQRE profile of the base game and $\delta \geq 0$ without loss of generality.

Proposition 14. *For given observed strategy frequencies $(p + \delta, p - \delta)$, the sign of $\frac{\hat{\pi}_1^{EC} + \hat{\pi}_2^{EC}}{2} - \frac{\hat{\pi}_1^{EP} + \hat{\pi}_2^{EP}}{2}$ is the same as the sign of $c_{UL} + c_{DR}$.*

Proof. From the proof of Proposition 12, we know that given δ the profile $\hat{\pi}^{EP}(\delta)$ satisfies

$$\frac{p + \delta - \hat{\pi}_1^{EP}}{p - \delta - \hat{\pi}_2^{EP}} = -\frac{\Delta_2}{\Delta_1}. \tag{13}$$

Consider the case $c_{UL} + c_{DR} > 0$. Then $\Delta_2 > \Delta_1$, which, applying to (13), implies $p + \delta - \hat{\pi}_1^{EP} > -p + \delta + \hat{\pi}_2^{EP}$, which reduces to $\hat{\pi}_1^{EP} + \hat{\pi}_2^{EP} < 2p$. The cases $c_{UL} + c_{DR} < 0$ and $c_{UL} + c_{DR} = 0$ are analogous. Recalling that $\hat{\pi}_1^{EC} = \hat{\pi}_2^{EC} = p$ completes the argument. $\square$

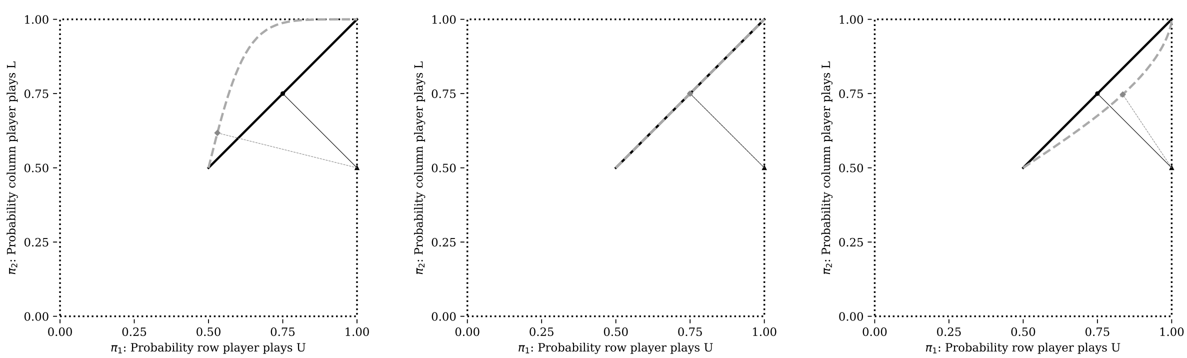

(a) $c_{UL} = 2, c_{DR} = 1$.   (b) $c_{UL} = 2, c_{DR} = -2$.   (c) $c_{UL} = 2, c_{DR} = -4$.

Figure 7: Examples of the three cases for the family of games in Table 5. The empirical strategy frequencies illustrated are $\left(1, \frac{1}{2}\right)$ for all three examples, all of which are mapped to $\left(\frac{3}{4}, \frac{3}{4}\right)$ by the equilibrium correspondence approach.

The three cases in Proposition 14 are illustrated in Figure 7. The game in Figure 7a is a Pareto coordination game; a qualitatively similar picture would arise if D and R were dominated strategies but the outcome (D, R) is not too bad ($c_{DR} < 0$ but $c_{UL} > -c_{DR}$). The game in Figure 7c is a game in which D and R are dominated strategies, and the outcome arising from the combination of D and R is particularly bad for both players. The knife-edge case in Figure 7b arises in, for example, a two-player linear voluntary contributions game. In that case, the empirical payoff approach always gets Δ_1 and Δ_2 correct, because those differences are independent of the strategies played.

Figure 7 alone is not enough to determine whether there are systematic differences in λ between the two approaches. Figure 8 plots level sets for $\hat{\lambda}^{EC}$ and $\hat{\lambda}^{EP}$. Compared to the Ochs game, we observe the useful property that in this game, the level sets of $\hat{\lambda}^{EC}$ and $\hat{\lambda}^{EP}$ are tangent, for given λ, at the point they cross the locus of LQRE profiles. We also see that for these example games,

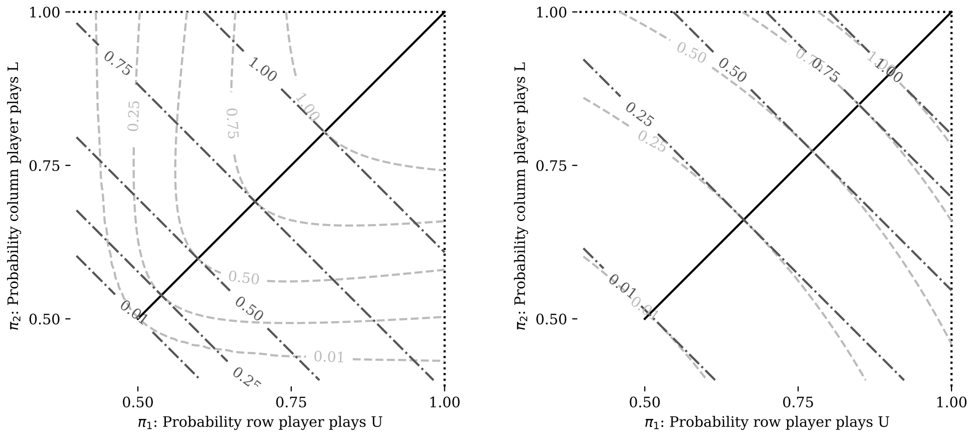

(a) $c_{UL} = 2, c_{DR} = 1$. (b) $c_{UL} = 2, c_{DR} = -4$.

Figure 8: Level sets of $\hat{\lambda}^{EC}$ and $\hat{\lambda}^{EP}$ for games from Figure 7 with symmetric LQRE profiles. The solid line is the locus of LQRE profiles for the base game. The darker dot-dashed lines correspond to level sets of $\hat{\lambda}^{EC}$, and lighter dashed lines to level sets of $\hat{\lambda}^{EP}$.

the curvature of the level sets of $\hat{\lambda}^{EP}$ is game-specific. We now formalize this observation to characterize the cases when $\hat{\lambda}^{EC} \gtrless \hat{\lambda}^{EP}$.

Lemma 15. *For the family of games in Table 5, the empirical payoff approach profile $\hat{\pi}^{EP}$ satisfies $\Delta_2 \hat{\pi}_2^{EP}(1-\hat{\pi}_2^{EP}) - \Delta_1 \hat{\pi}_1^{EP}(1-\hat{\pi}_1^{EP}) > 0$.*

Proof. Given realized strategy frequencies p, the level sets of the log-likelihood function satisfy

$$p_1 \ln \pi_1 + (1-p_1)\ln(1-\pi_1) + p_2 \ln \pi_2 + (1-p_2)\ln(1-\pi_2) = K \tag{14}$$

for any given K. Because in the auxiliary game π_1 and π_2 are both increasing in λ we can parameterize the locus of LQRE profiles by $(\pi_1, \pi_2(\pi_1))$. Differentiating (14) with respect to π_1 and rearranging gives

$$\frac{\pi_1 - p_1}{\pi_2 - p_2} = -\frac{\pi_1(1-\pi_1)}{\pi_2(1-\pi_2)}\pi_2'. \tag{15}$$

Geometrically, this says that along any line segment starting at p, the level sets of the likelihood function are orthogonal to that line segment. Evaluating (15) at $\hat{\pi}^{EP}$ and applying Proposition 12,

$$\begin{aligned} \frac{\Delta_2}{\Delta_1} &= \frac{\hat{\pi}_1^{EP}(1-\hat{\pi}_1^{EP})}{\hat{\pi}_2^{EP}(1-\hat{\pi}_2^{EP})}\pi_2'(\hat{\pi}_1^{EP}) \\ \pi_2'(\hat{\pi}_1^{EP}) &= \frac{\Delta_2 \hat{\pi}_2^{EP}(1-\hat{\pi}_2^{EP})}{\Delta_1 \hat{\pi}_1^{EP}(1-\hat{\pi}_1^{EP})} > 0, \end{aligned}$$

because $\pi_2' > 0$ due to the monotonicity of the locus of LQRE profiles. As a geometric interpretation, note that $\pi_2(\lambda) - \pi_1(\lambda)$ is quasiconcave in λ; therefore, this result states that the empirical payoff estimator $\hat{\lambda}$ occurs at a lower value of λ than the one that maximises the difference between π_2 and π_1. □

Proposition 16. *For the family of games in Table 5, the sign of* $\hat{\lambda}^{EC}(p) - \hat{\lambda}^{EP}(p)$ *is the same as the sign of* $c_{UL} + c_{DR}$.

Proof. Again from the proof of Proposition 12, we know that given δ,

$$\Delta_1(p + \delta - \hat{\pi}_1^{EP}) + \Delta_2(p - \delta - \hat{\pi}_2^{EP}) = 0. \tag{16}$$

We can rearrange this to

$$\begin{aligned}\Delta_1\hat{\pi}_1^{EP} + \Delta_2\hat{\pi}_2^{EP} &= (\Delta_1 - \Delta_2)\delta + (\Delta_1 + \Delta_2)p \\ &= -2\kappa\delta^2 + (\Delta_1 + \Delta_2)p.\end{aligned}$$

Let $\phi(x) = \frac{e^x}{1+e^x}$, and note that $\phi'(x) = \phi(x)(1 - \phi(x))$. Because the $\hat{\pi}_i^{EP}$ are an LQRE profile of the auxiliary game for some $\hat{\lambda}^{EP}$,

$$\Delta_1\phi(\hat{\lambda}^{EP}\Delta_1) + \Delta_2\phi(\hat{\lambda}^{EP}\Delta_2) = -2\kappa\delta^2 + (\Delta_1 + \Delta_2)p. \tag{17}$$

Differentiating with respect to δ, noting that $\hat{\lambda}^{EP}$, and Δ_i are all functions of δ,

$$\begin{aligned}\Delta_1'\phi(\hat{\lambda}^{EP}\Delta_1) + \Delta_1\phi'(\hat{\lambda}^{EP}\Delta_1)(\hat{\lambda}^{EP'}\Delta_1 + \hat{\lambda}^{EP}\Delta_1') & \\ +\Delta_2'\phi(\hat{\lambda}^{EP}\Delta_2) + \Delta_2\phi'(\hat{\lambda}^{EP}\Delta_2)(\hat{\lambda}^{EP'}\Delta_2 + \hat{\lambda}^{EP}\Delta_2') &= -4\kappa\delta + (\Delta_1' + \Delta_2')p\end{aligned}$$

Let $\kappa = c_{UL} + c_{DR}$ and note that $\Delta_1' = -\kappa$ and $\Delta_2' = \kappa$,

$$\begin{aligned}-\kappa\phi(\hat{\lambda}^{EP}\Delta_1) + \Delta_1\phi'(\hat{\lambda}^{EP}\Delta_1)(\hat{\lambda}^{EP'}\Delta_1 - \hat{\lambda}^{EP}\kappa) & \\ +\kappa\phi(\hat{\lambda}^{EP}\Delta_2) + \Delta_2\phi'(\hat{\lambda}^{EP}\Delta_2)(\hat{\lambda}^{EP'}\Delta_2 + \hat{\lambda}^{EP}\kappa) &= -4\kappa\delta \\ \hat{\lambda}^{EP'}\left[\Delta_1^2\phi'(\hat{\lambda}^{EP}\Delta_1) + \Delta_2^2\phi'(\hat{\lambda}^{EP}\Delta_2)\right] & \\ +\kappa\left\{\left[\phi(\hat{\lambda}^{EP}\Delta_2) - \phi(\hat{\lambda}^{EP}\Delta_1)\right] + \hat{\lambda}^{EP}\left[\Delta_2\phi'(\hat{\lambda}^{EP}\Delta_2) - \Delta_1\phi'(\hat{\lambda}^{EP}\Delta_1)\right]\right\} &= -4\kappa\delta\end{aligned}$$

Consider the case $\kappa > 0$; in this case we want to show that $\hat{\lambda}^{EP'} < 0$. The expression in square brackets multiplying $\hat{\lambda}^{EP'} < 0$ is positive, and the right-hand side is negative. To complete the argument the expression in curly braces multiplying κ must be positive. $\phi(\lambda\Delta_2) > \phi(\lambda\Delta_1)$ for all $\lambda > 0$ because $\Delta_2 > \Delta_1$, and $\Delta_2\phi'(\hat{\lambda}^{EP}\Delta_2) - \Delta_1\phi'(\hat{\lambda}^{EP}\Delta_1) > 0$ from Lemma 15. The argument for $\kappa < 0$ is analogous.

The result then follows from observing that when $\delta = 0$, $\hat{\lambda}^{EC} = \hat{\lambda}^{EP}$, and $\hat{\lambda}^{EC}$ is constant with respect to δ given p. □

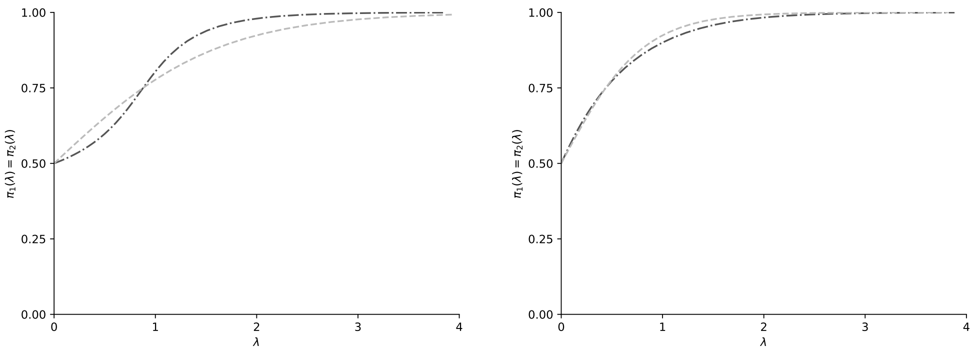

(a) $c_{UL} = 2, c_{DR} = 1$. (b) $c_{UL} = 2, c_{DR} = -4$.

Figure 9: Correspondence plots of $\pi_1^{EC}(\lambda) = \pi_2^{EC}(\lambda)$ and $\pi_1^{EP}(\lambda) = \pi_2^{EP}(\lambda)$ for games from Figure 7 with symmetric LQRE profiles, given observed strategy frequencies $p_1 = p_2 = \frac{3}{4}$. The darker dot-dashed lines correspond to the correspondence π^{EC}, and lighter dashed lines to π^{EP}.

Some intuition for these results can be gained by recalling the difference between the locus of LQRE profiles and the LQRE correspondence. In this family of games, when the realized strategy frequencies are the same, $p_1 = p_2$, the locus of LQRE profiles of the auxiliary game is exactly the same as that of the base game. However, the mapping from λ to profiles is not the same. Figure 9 demonstrates this for the case when the observed frequencies are $p_1 = p_2 = \frac{3}{4}$. In the auxiliary game, the value of $p_1 = p_2$ sets the strength of the incentives between U and D (respectively, L and R), and this is constant for all λ. Consider the case of the Pareto-coordination game as in Figure 9a. For small λ, this overstates the strategic incentive, relative to that of the LQRE of the base game, and so π^{EP} increases more rapidly. For large λ the situation is reversed; the incentives are stronger in the base game and so the probability of U (respectively L) increases more rapidly. This is reversed when c_{DR} is very low; this situation is depicted in Figure 9b.

This example shows that the patterns of bias in $\hat{\lambda}^{EP}$ are not coming from the fact that randomization is important in the asymmetric and generalized matching pennies games. In cases in which a limiting pure-strategy equilibrium is relevant, increasing the probability a strategy is played by one player often sharpens the incentives of another player, such as occurs in games with a coordination aspect. Omitting this reinforcement effect in LQRE usually leads to a systematic bias in $\hat{\lambda}^{EP}$, which can be in either direction depending on the structure of the game.

5 Feasibility of the equilibrium correspondence approach

In many instances in which we turn to numerical calculation to solve optimization problems, we do so because we cannot easily characterize the optimizer because the objective function is complex, while the constraints on the feasible set are relatively simple. This is the use case for which optimization routines in most software packages are targeted. The optimization problem defined in (5) is interestingly different because the objective function is very easy to work with - the log-likelihood is strictly concave – but the structure of the feasible set $\mathcal{P}(\Gamma)$ is potentially complex for a non-trivial game.

Because $(0, \pi^0)$ is always path-connected to at least one $(\infty, \pi^\star)$, where $\pi^\star$ is a Nash equilibrium, computing the full feasible set in (5) is a method for computing a Nash equilibrium. The problem of computing a Nash equilibrium of a finite game has been shown to be PPAD-complete (Daskalakis et al., 2009). This is bad news, because it is believed that problems in PPAD are hard, and no algorithms are known to solve these problems in time polynomial in the size of the input. So, even computing the principal branch of the LQRE correspondence may be difficult. Meanwhile, the problem of determining whether a game has a second Nash equilibrium is NP-complete, so optimizing over all LQRE, including those not on the principal branch, is a difficult problem.

However, these complexity class results apply to *all* problems, that is, all possible games. Further, even if it is the case that any algorithm must take time that grows exponentially in the size of the game, nevertheless exponential growth may be sufficiently slow as not to matter significantly when working with games of even moderate size. We observe that the practical problem of taking LQRE to data is generally not done on randomly-generated games. In experimental settings, experimenters usually choose games which have some structure on their payoff functions, and the (effective) size of the strategy space in most experimental games is not too large.[10] So we turn to a discussion which outlines how computer codes which solve (5) operate, and the practical computational load associated with those algorithms as a function of the types of games most frequently studied. The methods available in the Gambit package (McKelvey et al., 2022) for computing LQRE and doing maximum-likelihood estimation are based on the methods below. As part of an electronic appendix to this paper we provide standalone implementations in pure Python to complement the exposition.

The problem of computing points along a smooth curve is well-suited to the use of the *predictor-corrector* (PC) method. We will develop the basic ideas briefly, based on the treatment of Allgower and Georg (2003). Following the developments in Turocy (2005) and Turocy (2010), we can write a curve $m \in \mathcal{M}$ as $m(s) = (\rho(s), \lambda(s))$ satisfying $H(m(s)) = 0$, where ρ is the vector of log-probabilities, $\rho_{ij}(s) = \log \pi_{ij}(s)$, and the curve m is parameterized by arclength s. We can then

[10]In referring to the "effective" size of a strategy space, we are observing that, while an experiment may permit participants to select a number from, say, 0 to 1000 inclusive, in practice choices will be concentrated on "accessible" numbers such as multiples of 10 or 100, and the vast majority of strategies will not appear at all in the experiment's dataset.

re-state the LQRE conditions (2) as

$$\begin{array}{ll} H_{ij} : \rho_{i,j+1} - \rho_{ij} - \lambda(u_{i,j+1}(\mathrm{e}^{\rho}) - u_{ij}(\mathrm{e}^{\rho})) & \forall i,j : 1 \le i \le n, 1 \le j \le J_i - 1 \\ H_{i\Sigma} : \sum_{j=1}^{J_i} \mathrm{e}^{\rho}_{ij} - 1 & \forall i : i \le i \le n. \end{array} \tag{18}$$

The smoothness of the curve justifies treating the problem as a differential equation on the parameter s, and so $H'(m(s))\nabla m(s) = 0$, where $||\nabla m(s)|| = 1$ is chosen by convention. This equation allows us to compute the direction of the curve ∇m at a given point $m(s)$ based on the system's Jacobian $H'(m(s))$ at that point.

The PC method[11] produces a sequence $\{m_t\}$ of points along the curve. The method requires a starting point m_0 known to be on the curve; for example, starting from $(0, \pi^0)$ allows the principal branch to be traced. Given a point m_t, a *predictor* for m_{t+1} is determined by $\hat{m}^0_{t+1} = m_t + h\nabla(m_t)$; that is, a step of length h is taken in the direction suggested by the tangent to the Jacobian at m_t. This predictor step is typical for any numerical differential-equation solver. More can be done, however, by taking advantage of the knowledge that $H(m(s)) = 0$; for example, we can apply Newton's method for finding a zero of a system of equations, an iterative process which converges quadratically to a solution. Specifically, a sequence of correctors can be computed via $\hat{m}^{k+1}_{t+1} = \hat{m}^k_{t+1} - H'(\hat{m}^k_{t+1})^+ H(\hat{m}^k_{t+1})$, where A^+ denotes the Moore-Penrose inverse of A; this iteration can be done until a required accuracy $||\hat{m}^k_{t+1}|| < \varepsilon$ is obtained. The point accepted under this criterion then becomes m_{t+1}.

We can contrast the PC method with the approach taken in the sample MATLAB codes in Goeree et al. (2016). In those codes, the equations defining logit responses are transformed into a minimization problem by taking the sum of squares of "belief errors," that is, the differences between the belief probabilities and the logit responses to those beliefs. It is true that the global minima of this function correspond exactly to LQRE. However, as pointed out by Judd (1998), the technique of converting finding the zeroes of a system of equations to one in which a function is minimized has a number of unattractive numerical properties. In particular, the condition number of the Hessian of the objective function is roughly the square of the condition number of the Jacobian of the system of equations. As a result, numerical optimization routines may struggle to find a solution even when methods for solving the system of equations directly experience no problems.[12]

This brief discussion of the ideas behind predictor-corrector methods is useful reference for understanding the practical computational expectations for being able to solve problem (5). First,

[11] Properly speaking, there are many variations on the predictor-corrector idea for numerical continuation, and as such referring to "a" PC method would be more proper. In this paper we will discuss one implementation, and as such use the definite article where it is more natural and where reference to a specific implementation is unambiguous from the context.

[12] See the discussion on pages 171–173 of Judd (1998) for further details.

we note that the PC method computes ∇m_t for every point m_t. The values of m_t and ∇m_t are exactly the data required to be able to evaluate the first-order condition (5). Newton's method can be adapted to find a zero of a function such as (5) with desired precision and likewise benefits from quadratic convergence. This formalizes the observation that all of the hard computational work for solving problem (5) is in computing points on or sufficiently close to the curves in $\mathcal{M}$; having done this computation, finding zeroes of the first-order condition can be done at essentially zero extra computational effort. We will therefore focus our discussion on practical aspects of the process of tracing out a given curve m.

So far, we have not discussed the choice of the hyperparameter h of the PC method. Clearly h cannot be chosen as a constant, because the curves comprising $\mathcal{L}(\Gamma)$ which limit to a Nash equilibrium have an infinitely-long arclength.[13] PC methods generally use an adaptive steplength. The idea behind steplength adaptation is that the predictor step just needs to be "good enough" that the corrector step converges quickly to a point on the curve. Given a current steplength h and a point $m(s)$ on the curve, if $m(s) + h\nabla m(s)$ turns out to be a point on or close to the curve, then heuristically it would make sense to try a larger steplength on the next iteration. Conversely, if convergence at the predicted point is slow, or fails entirely, then trying again with a shorter steplength would be indicated.

We can illustrate some intuition for the interaction of the predictor and corrector steps, and steplength adaptation strategies, by looking at the PC method as applied to auxiliary games.[14] Because we know that there are no turning points in an auxiliary game, we can simplify the exposition by taking λ as the parameter. For each player i and all strategies j and k of player i, we have that

$$\begin{aligned} \rho_{ij}(\lambda) - \rho_{ik}(\lambda) &= \lambda(u_{ij} - u_{ik}) \\ \rho'_{ij}(\lambda) - \rho'_{ik}(\lambda) &= u_{ij} - u_{ik} \end{aligned}$$

[13] The arclength can be made finite by instead re-parameterizing the curve via a transformation similar to $\nu \equiv \frac{\lambda}{1+\lambda}$, referred to humorously in Turocy (2005) as the "Texas" parameterization. Even with this parameterization a fixed steplength is not satisfying, as the transformation rescales the regions of the noise/precision parameter on which the behavior of the LQRE set may be the most interesting, and so an adaptive steplength would still be indicated.

[14] Although numerical continuation is not necessary in these cases, as the LQRE can be found through direct computation of the formulas, these are useful for building understanding of how PC methods and steplength adaptation work.

We can differentiate the requirement that probabilities sum to one for each player to obtain

$$\begin{aligned}
\sum_{j=1}^{J_i} \rho'_{ij}(\lambda) \mathrm{e}^{\rho_{ij}(\lambda)} &= 0 \\
\sum_{j=1}^{J_i} \left[\rho'_{ik}(\lambda) + (u_{ij} - u_{ik})\right] \mathrm{e}^{\rho_{ij}(\lambda)} &= 0 \\
\rho'_{ik}(\lambda) + \sum_{j=1}^{J_i} (u_{ij} - u_{ik}) \mathrm{e}^{\rho_{ij}(\lambda)} &= 0 \\
\rho'_{ik}(\lambda) &= -\sum_{j=1}^{J_i} (u_{ij} - u_{ik}) \pi_{ij}(\lambda) \qquad (19)
\end{aligned}$$

The significance of (19) is that, given $(\lambda, \rho(\lambda))$, it is used to predict the next point $(\lambda+h, \hat{\rho}(\lambda+h))$. We now observe that this predictor is quite good, in that the predictor always satisfies the relative-payoff equation exactly:

$$\begin{aligned}
\hat{\rho}_{ij} - \hat{\rho}_{ik} &= \left[\rho_{ij} - h \sum_{l=1}^{J_i} (u_{il} - u_{ij}) \pi_{il}\right] - \left[\rho_{ik} - h \sum_{l=1}^{J_i} (u_{il} - u_{ik}) \pi_{il}\right] \\
&= \rho_{ij} - \rho_{ik} + h \left[\sum_{l=1}^{J_i} (u_{ij} - u_{ik}) \pi_{il}\right] \\
&= \lambda (u_{ij} - u_{ik}) + h (u_{ij} - u_{ik}) \\
&= (\lambda + h)(u_{ij} - u_{ik}).
\end{aligned}$$

The predictor is not perfect because the system is not linear in log-probabilities; the sum-to-one equations for each player do not balance exactly when using this predictor. The Newton step in the corrector recognizes this, and takes a step that re-scales the total probabilities while keeping the relative log-probabilities in their correct proportions. In fact the corrector phase needs only one step to find a point on the curve, as the step will always complete the normalization exactly on the first iteration. Because the convergence is so rapid, this is evidence that the steplength can be increased, and so traversal of the curve can proceed using relatively few steps. The PC method is therefore surprisingly competitive for computing the LQRE of auxiliary games when compared to solving the equations directly, as it does more or less the same computational operations.

However, because the problem of computing a Nash equilibrium is PPAD-complete, we know that there must be games for which the PC method will not work as elegantly as this. The running time of the PC method is determined by the number of steps it must take, and the computational time involved in taking one step. We consider these in turn.

At a given point $m(s)$, the tangent $\nabla m(s)$ used to form the predictor incorporates two effects, which we can think of as the **payoff effect** and the **strategic effect**.[15] The payoff effect is determined by applying the replicator dynamics, which as we have just seen, is well-predicted by a linear approximation in log-probabilities. The strategic effect comes from the adaptation of each player's play to the anticipated play of others. The predictor step operates, roughly, by assuming that the strategic effect at $m(s)$ is a good approximation to the effect at $m(s+h)$. The extent to which this is true will influence the steplength adaptation, and therefore the number of steps needed; this is different across games, and even at different parts of the LQRE correspondence for a given game.

Irrespective of how the number of steps required by the PC method depends on the features of a game, an important practical determinant of running times is the time it takes to complete a step. It is important to remember that for practical application, it is not necessarily the case that algorithms with running times that scale exponentially in the size of their input are necessarily impractical and those which scale polynomially are practical; an exponential algorithm with running times growing slowly over the range of problem sizes of interest may beat a polynomial one with running times growing more rapidly over the range. The PC method requires information both about the function $H(m(s))$ defining the curve, and its Jacobian $H'(m(s))$, at each of the points $\hat{m}^k t$ considered as a predictor or corrector. The method then does linear algebra to compute the relevant tangent for the predictor step, and steps for the correctors. These operations are done on matrices of size $J \times J$ and $J \times (J+1)$, so the total number of strategies J is the determinant of the complexity of this part of the algorithm.

However, experience with implementing the PC method in Gambit (McKelvey et al., 2022) has shown that it is neither the number of steps nor the linear algebra at each step which dominates the running time. Instead, the overwhelming majority of computational time is spent computing expected payoffs. Here, one of the attractive features of LQRE, that it has full support for all finite λ, implies a curse of dimensionality. We must visit all $\prod_{i=1}^n J_i$ contingencies of the game to compute the expected payoffs of all of player i's strategies, meaning that the cost of computing $H(m)$ is proportional to $N \prod_{i=1}^n J_i$.

However, in the most generic implementations, it is the Jacobian which consumes most of the computational cycles. For each of the H_{ij} equations, computing the Jacobian precisely implies:

1. Computing $\frac{\partial H_{ij}}{\partial \lambda}$.
2. Computing $\frac{\partial H_{ij}}{\partial \rho_{lm}}$ for every player $l \neq i$ and every strategy m of that player l;

These two items correspond precisely to the payoff effect and the strategic effect, respectively.

[15] This builds off of the observations in Turocy (2005) relating LQRE tracing to the Tracing Procedure of Harsanyi and Selten (1988).

The column of the Jacobian consisting of $\frac{\partial H_{ij}}{\partial \lambda}$ is exactly the replicator dynamics, evaluated at the current vector of strategy payoffs. The columns with the derivatives with respect to strategy probabilities capture the changes in strategy values. Here we can exactly pinpoint one of the ways in which the empirical payoff approach saves computational time: it sets all of those cross-partials with respect to strategies to zero. In the simplest implementation of expected payoff calculations, this saves visits to $\sum_{i=1}^{n} \sum_{k \neq i} (J_i + J_k) \prod_{l \neq i,k} J_l$ contingencies in the game[16] for each step of the PC method.

This analysis indicates that the feasibility in practice of using the equilibrium correspondence approach is closely tied to the computation of the Jacobian. We now discuss some techniques which can be used to improve the performance of estimation, in some cases by many orders of magnitude.

5.1 Take advantage of the structure of the game

The straightforward way to compute expected payoffs is to visit all the contingencies of the game. This, for example, is what (as of this writing) Gambit does in its generic implementation of maximum likelihood estimation for LQRE. This is always correct, and for some games it is the best that can be done.

However, many of the games we actually study, and especially those which are subjected to experimental or empirical analysis, are ones which have some structure in how payoffs are determined. In applying LQRE, numerical methods are essential because it is only in rare instances where closed-form expressions for LQRE probabilities can be obtained. These examples can be linked to cases in which there are substantial efficiency gains to be had in computing the Jacobian.

One example is the all-pay auction, analysed by Anderson et al. (1998a) in a continuous-bid format. A property of the all-pay auction that permitted this analysis is that the difference in expected payoffs between two adjacent bids in the strategy space is very easy to compute. To fix ideas, suppose the set of feasible bids is a discrete set $\{b_0, b_1, \ldots, b_J\}$. In the two-bidder case, the derivatives of the equation H_{ij} relating the probability of bidding b_j to that of bidding b_{j+1} are all zero, except for the columns corresponding to other bidder(s) bidding those bids. A custom implementation of the Jacobian can thus get enormous speedups. Other games, such as the "11-20 game" of Arad and Rubinstein (2012), have a similar structure. We observe that a similar observation applies to first-price auctions. First-price auctions typically are modeled as Bayesian games; the predictor-corrector approach described above applies equally to the agent LQRE of McKelvey and Palfrey (1998), and the resulting Jacobian will be sparse for the same reasons.

Anderson et al. (1998b) provided an analysis of voluntary contributions public goods games

[16] Where the product over an empty range is defined by convention to be 1.

using LQRE. In the linear case, in which the cost of contribution is linear in the size of contribution, and the benefit from contribution is constant, then, as noted previously, the cross-partial derivatives in the Jacobian are constant. This arises from the linear separability of benefits and costs. Therefore, they are also able to analyse the case of quadratic, but still separable, costs.

The benefits of custom routines for computing expected payoffs are not always so dramatic. In computing, it frequently can make practical sense to trade off computational time for the time of people writing custom code, as the implied cost of the latter is high. However, the Jacobian is the critical path in a PC method, and so investing in more efficient implementation therefore is often worth it. Computational libraries in other domains increasingly do introspection in order to optimise an execution plan (for example, database engines do this to optimise query plans). Similar techniques may be applicable to automate the discovery of ways to compute expected payoffs in a given family of games with a minimum of computational operations.

5.2 Approximate the Jacobian

It may not always be practical to analyze a game to determine the optimal approach for computing the Jacobian. Further, in some games, there might be little opportunity to do so. Consider for example a second-price auction with private information and correlated types or values. Unlike the first-price and all-pay auctions, in the second-price auction the expected payoff to a bid depends on the probabilities of all of the bids below it, because conditional on winning any of those lower bids could be the one that sets the price the winner pays. The task is further complicated by correlation among types and any common value component, meaning that inferring the value of the object conditional on winning, given that all types follow a strategy that has full support over all possible bids, is complicated. Clever computational tricks in the Jacobian will have only limited benefit in such settings, and custom code will be prone to errors which are difficult to detect.

There are some interesting observations which arise from consideration of the role of the Jacobian in a PC method. At the predictor step, $H'(m)$ is important because it is the information required to compute ∇m, which determines the direction of the step. However, because there is a subsequent corrector step, all that is important about the predictor step is that it is "good enough" that the corrector step is then able to converge. Further, the corrector step also does not require an accurate Jacobian. The quadratic convergence rate of a Newton-type iteration does assume access to the Jacobian at each iteration. However, less accurate estimates of the Jacobian in general only slow convergence, but do not destroy it; the secant method for finding a root, for example, requires only $H(m)$, not $H'(m)$, at the cost of convergence taking more iterations.[17]

[17] In fact, the PC method implemented in the ancillary materials for this paper computes the Jacobian at the predictor, and uses that resulting Jacobian as an approximation to the true Jacobian for the corrector iterations.

As an alternative, there are variations on the PC method which use the information obtained from evaluations of $H(m)$ to maintain an approximation to the Jacobian. One such implementation is given by Allgower and Georg (2003). The basic idea is that, given an initial value of the Jacobian at some point, that value can be updated using the information already computed by taking a predictor step. This update is only an approximation, as it uses just a fraction of the information that is contained in the full Jacobian. In the electronic auxiliary materials with this paper we include a sample implementation in Python of such a method, based on the code of Allgower and Georg (2003).[18]

There is an inherent tradeoff between using a Jacobian-approximation method versus one in which it is computed exactly. The more precise the Jacobian is, the more accurate the predictor step will be, and therefore longer steps can be taken. With an approximated or estimated Jacobian, shorter steps will be taken, but each step will be computationally less expensive.

5.3 Approximation methods

As we have already observed, computing the Jacobian requires computing $\frac{\partial u_{ij}}{\partial \pi_{lm}}$, for all pairs of players $i, l \in N$ and all pairs of their strategies $j \in S_i$ and $m \in S_l$. In the case of two player games in normal form, this is a constant. Therefore, the only variable part of the Jacobian is the column corresponding to $\nabla_\lambda H(m)$. Each of these entries is just the difference of the expected payoffs of two strategies, and those expected payoffs are already required for computing $H(m)$. It follows that for two-player games, no extra computational work is required to provide an exact Jacobian at every step. This observation extends to computing the Jacobian of agent LQRE for Bayesian games. This observation suggests that for two-player games, computational feasibility is never, or at most very rarely, a justification for the empirical payoff method, and the equilibrium correspondence approach should *always* be used.

This special property of the Jacobian is also enjoyed by *polymatrix games* (Janovskaya, 1968; Howson, 1972). In a polymatrix game, the payoff to player i from choosing strategy j against a given pure strategy profile s can be written in the form

$$u_{ij}(s) = \sum_{l \neq i} v^{il}(j, s_k). \tag{20}$$

In other words, the game can be expressed as a set of bimatrix games between each pair of players i and l, with payoff function v^{il}.

Approximation by polymatrix games has been used to good effect in computing Nash equilib-

[18]There are definitely errors in the code published in Allgower and Georg (2003). We believe we have corrected at least some of these in the sample codes, but the code does come with the disclaimer that it is for illustrative purposes only.

ria. Govindan and Wilson (2003) introduced a path-following global Newton method to compute equilibria. Briefly, the idea of the method is to perturb the payoffs of the game such that there exists a unique and easy-to-find equilibrium, such as one in which all players have a strictly dominant strategy. Then, the structure theorem for Nash equilibrium can be exploited to trace along the set of equilibria back to the original game of interest. This method is intuitively appealing, but comes with some implementation subtleties and can be quite slow. Govindan and Wilson (2004) subsequently introduced a complementary approach, which approximates the game by a sequence of polymatrix games; they find that when method based on polymatrix approximations converges, it typically finds equilibria much faster than the global Newton method.

These considerations suggest a computational strategy for games with more than two players, which can incorporate some of the ideas in the empirical payoff approach without throwing out entirely the strategic information of a game. If Jacobian-approximation methods are not suitable because the updating of the approximation is too poor to permit step sizes to be increased, one can approximate the game by treating it as a polymatrix game, in which, following the idea of Govindan and Wilson (2004), the payoff matrices for each pair of players are determined by taking expected values given the empirical strategy frequencies of the other $n-2$ players. This method retains the computational advantages of two-player games - that computing expected payoffs can be organised quite efficiently by row and column multiplications, and that the Jacobian matrices require only filling in one column based on those expected payoffs - while also incorporating a first-order approximation to the strategic aspects.

6 Discussion

We have extended the discussion of the empirical payoff approach to estimating LQRE presented in Goeree et al. (2016) by providing the first systematic analysis of the approach's properties. We have focused principally on 2×2 games. Although the empirical payoff approach is not necessary computationally for these simple games, they are nevertheless rich enough to illustrate the systematic differences between the empirical payoff approach and the equilibrium correspondence approach. We have isolated and illustrated several important differences between the two approaches.

1. The results of the empirical payoff approach are sensitive to small movements in observed data in regions of the space of mixed strategy profiles where the ordering of strategies by their expected payoffs change. This is in particular a concern in games in which active randomization is behaviorally relevant, including when there are equilibria in which randomized strategies are played.

2. The comparative statics of the strategy profiles, and values of the precision parameter λ, as estimated by the empirical payoff approach, may be counterintuitive. Along a sequence of strategy frequencies which equilibrium analysis suggests is increasing in precision, the empirical payoff approach can actually claim behavior is becoming *less* precise. This arises in part because the LQRE correspondence used by the empirical payoff approach is generally not a good approximation to the correspondence of the original game.

3. Even in games in which the locus of LQRE of the original game is sufficiently simple that the empirical payoff approach's approximation is not too bad, there remain systematic biases in the estimates of the precision parameter λ. The direction of the bias is game specific: although in most games the tendency is for the empirical payoff approach to underestimate the precision of play, counterexamples can be constructed where the systematic bias is to overestimate precision.

The root of all of these differences is the way the empirical payoff approach approximates the incentives in the game. Compared to the information available in the game, the empirical payoff approach uses just one small piece of information: the expected payoffs associated with each strategy against the sample of play. The implications of using such coarse information, and discarding the information carried by the original game, is illustrated by the structure of the Jacobian of the system of equations defining the LQRE correspondence: the entries associated with the strategic structure are set to zero, leaving only the replicator dynamics component of LQRE.

Our results show that even for moderate-sized datasets, the empirical payoff approach does not produce particularly satisfying results. Perhaps the most important achievement of McKelvey and Palfrey (1995) in defining the quantal response equilibrium is that it incorporates *both* quantal response *and* equilibrium. The empirical payoff approach discards the latter. To say that one "estimates QRE using the empirical payoff approach" therefore borders on an oxymoron; what is returned by the empirical payoff approach is not necessarily even approximately an equilibrium of the base game.

Nevertheless, the motivation of wanting to approximate QRE that underpins the empirical payoff approach does have valid roots in computational complexity. The empirical payoff approach merely goes too far, discarding too much important information about the game. Practical analysis of predictor-corrector methods for computing points on the LQRE correspondence sheds light on where improvements can be made. The empirical payoff approach approximates the Jacobian of the system by zeroing out most elements. Instead, approaches based on various heuristic approximations would seem promising, allowing at least some of the strategic information of the game to be incorporated properly into estimations. It is also worth noting that if the objective of a tracing of the correspondence via numerical continuation is solving the maximum likelihood problem, algorithms can use a greater tolerance for the precision of computing exactly the zeroes of the equations until a critical value of the likelihood function is being approached. The toolkit of numerical continuation - with proper tuning of the hyperparameters of algorithms - is a more promising approach to estimation and, combined with the modern scientific computing ecosystem, should render the empirical payoff approach in its current form obsolete.

We note in closing that our discussion of the computational aspects of tracing a branch of the LQRE correspondence is motivated by its practical relevance in taking LQRE to data as a structural model, and so we have not covered comprehensively all of the finer points involved in producing production-quality implementations. One issue which can catch out naive implementations of LQRE estimation code is bifurcations in the graph of LQRE. Goeree et al. (2016) mention this in the context of a battle-of-the-sexes game. Although the results of McKelvey and Palfrey (1995) show that generically bifurcations do not occur, in applications of game theory we do not choose the games we study randomly. As a result it is not particularly uncommon in applications for games with multiple equilibria to exhibit bifurcations. Goeree et al. (2016) assert (p. 153, footnote 9) that a "path-following algorithm ... generally runs into problems when the bifurcation value of λ is reached." We can briefly make this statement more precise, and offer a more optimistic view. Taking the battle-of-the-sexes example, the bifurcation point is where two curves in $\mathcal{L}(\Gamma)$ intersect. Looking at each curve in isolation, the tangent to the curve changes continuously when passing through that point. When the PC method passes through such a point in tracing a curve, the most common outcome is that it continues "straight through" the bifurcation. However, this event can easily be detected, because when passing through the bifurcation point, the sign of the tangent vector reverses.[19] The identification and numerical analysis of bifurcation points is a well-developed area;[20] Allgower and Georg (2003) provide a brief outline of some of the standard techniques for incorporating bifurcation analysis into PC methods.

[19]A common bug in home-grown PC code is that it does not check for the change in sign. Doing so results in the path following converging to, and getting stuck at, the bifurcation point. This behavior is a bug in the implementation, and not a problem inherent with numerical continuation.

[20]The study of bifurcation points is driven, for example, by the study of systems of equations corresponding to physical systems; in some of those applications, bifurcations can correspond to qualitative differences in the behavior of the system. The structure of bifurcations in general systems of equations can be quite complex; however, at least in the examples we are aware of, the bifurcations arising in LQRE correspondences are of relatively simple types.

References

Eugene L. Allgower and Kurt Georg. *Introduction to Numerical Continuation Methods*. Society for Industrial and Applied Mathematics, Society for Industrial and Applied Mathematics, 2003. doi: 10.1137/1.9780898719154. URL `https://epubs.siam.org/doi/abs/10.1137/1.9780898719154`.

Simon P. Anderson, Jacob K. Goeree, and Charles A. Holt. Rent seeking with bounded rationality: An analysis of the all-pay auction. *Journal of Political Economy*, 106:828–853, 1998a.

Simon P. Anderson, Jacob K. Goeree, and Charles A. Holt. A theoretical analysis of altruism and decision error in public goods games. *Journal of Public Economics*, 70:297–323, 1998b.

Ayala Arad and Ariel Rubinstein. The 11-20 money request game: A level-k reasoning study. *American Economic Review*, 102:3561–3573, 2012.

Patrick Bajari and Ali Hortacsu. Are structural estimates of auction models reasonable? Evidence from experimental data. *Journal of Political Economy*, 113:703–741, 2005.

Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39:195–259, 2009.

Evan Friedman and Felix Mauersberger. Quantal response equilibrium with symmetry: Representation and applications. Working paper dated 2022-11-17, 2022.

Jacob K. Goeree, Charles A. Holt, and Thomas R. Palfrey. Risk averse behavior in generalized matching pennies games. *Games and Economic Behavior*, 45:97–113, 2003.

Jacob K. Goeree, Charles A. Holt, and Thomas R. Palfrey. *Quantal Response Equilibrium: A Stochastic Theory of Games*. Princeton University Press, Princeton, NJ, 2016.

Srihari Govindan and Robert Wilson. A global Newton method to compute Nash equilibria. *Journal of Economic Theory*, 110(1):65–86, 2003.

Srihari Govindan and Robert Wilson. Computing Nash equilibria by iterated polymatrix approximation. *Journal of Economic Dynamics and Control*, 28(7):1229–1241, 2004.

John Harsanyi and Reinhard Selten. *A general theory of equilibrium selection in games*. MIT Press, Boston, MA, 1988.

John C. Harsanyi. Games with randomly distributed payoffs: A new rationale for mixed-strategy equilibrium points. *International Journal of Game Theory*, pages 1–23, 1973.

Joseph T. Howson. Equilibria of polymatrix games. *Management Science*, 18(5):312–318, 1972.

E. B. Janovskaya. Equilibrium points in polymatrix games. *Latvian Mathematical Collection*, 1968.

Kenneth L. Judd. *Numerical Methods in Economics*. MIT Press, Cambridge, MA, 1998.

Richard D. McKelvey and Thomas R. Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10:6–38, 1995.

Richard D. McKelvey and Thomas R. Palfrey. Quantal response equilibria for extensive form game. *Experimental Economics*, 1:9–41, 1998.

Richard D. McKelvey, Andrew M. McLennan, and Theodore L. Turocy. Gambit: Software tools for game theory. http://www.gambit-project.org, 2022.

Jack Ochs. Games with unique mixed strategy equilibria: An experimental study. *Games and Economic Behavior*, 10:174–189, 1995.

Reinhard Selten and Thorsten Chmura. Stationary concepts for experimental 2x2-games. *American Economic Review*, 98:938–966, 2008.

P. D. Taylor and L. Jonker. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 40:145–156, 1978.

Theodore L. Turocy. A dynamic homotopy interpretation of the logistic quantal response equilibrium correspondence. *Games and Economic Behavior*, 51:243–263, 2005.

Theodore L. Turocy. Computing sequential equilibria using agent quantal response equilibria. *Economic Theory*, 42:255–269, 2010.