

A Dynamic Homotopy Interpretation of the Logistic Quantal Response Equilibrium Correspondence

BY

BY THEODORE L. TUROCY

Department of Economics

Texas A&M University

College Station TX 77843

Phone: 979.862.8082

Fax: 979.847.8757

EMAIL: turocy@econ.tamu.edu

February 25, 2004

Abstract

This paper uses properties of the logistic quantal response equilibrium correspondence to compute Nash equilibria in finite games. It is shown that branches of the correspondence may be numerically traversed efficiently and securely. The method can be implemented on a multicomputer, allowing for application to large games. The path followed by the method has an interpretation analogous to Harsanyi and Selten's Tracing Procedure. As an application, it is shown that the principal branch of any quantal response equilibrium correspondence satisfying a monotonicity property converges to the risk-dominant equilibrium in 2×2 games.

JEL Classifications: C72, C88

Keywords: noncooperative games, computation of Nash equilibrium, quantal response, logit equilibrium.

1 Introduction

MCKELVEY AND PALFREY [15] introduced the concept of a quantal response equilibrium for finite n -person normal form games. This concept applies a quantal choice model, originated by LUCE [14], to interpret mixed strategy profiles as the observed distribution of strategy choices when the players observe payoffs with a random additive shock, and choose optimally according to those noisy payoffs. Subsequently, MCKELVEY AND PALFREY [16] extended the quantal response concept to extensive form games.

The additive disturbances associated with players' strategies may be drawn from any joint distribution satisfying an admissibility condition. In the case where these disturbances are drawn independently from the extreme value distribution with precision parameter λ , the form of the rule determining quantal response equilibrium choice probabilities is logistic. (See equation (1).) McKelvey and Palfrey refer to quantal response equilibria with this distribution of disturbances as logit equilibria, and this specification is widely used in analysis of subject behavior in laboratory games.

In the main body of the paper, the exposition will focus on quantal response equilibria in normal form games. The methods transfer to the agent logit quantal response equilibria of MCKELVEY AND PALFREY [16]. Appendix A shows how the methods can be applied equally to the agent specification.

The set of logit equilibria can be viewed as a correspondence from λ to the set of mixed strategy profiles. At $\lambda = 0$, this correspondence contains only the centroid. As λ approaches infinity, the correspondence converges to a (possibly strict) subset of the Nash equilibria of the game. Generically, the correspondence is structured such that there is a unique branch starting at the centroid at $\lambda = 0$, and limiting to a unique Nash equilibrium as λ approaches infinity.

McKelvey and Palfrey observe that this last fact may be used as the basis for an algorithm to compute a single Nash equilibrium of a game. They name this equilibrium the

logit solution. This proposed procedure makes use of the homotopy principle, as explained by GOVINDAN AND WILSON [6]:

Given a system of equations whose zeroes one wants to compute, first deform the system to one with a unique easily-computed solution, then reverse the deformation to trace (a selection of) solutions of the associated systems along the way to find a solution of the original system at the terminus.

As McKelvey and Palfrey noted, the existing procedure most similar in flavor to the tracing of logit equilibria is the Tracing Procedure of HARSANYI AND SELTEN [8].¹ HERINGS AND VAN DEN ELZEN [9] present a homotopy-based implementation of the Tracing Procedure, where the homotopy transforms the problem of playing a best reply to the initial beliefs to the problem of playing a best reply to opponents' actual play to form an equilibrium.

Homotopy approaches are common among algorithms for computing Nash equilibria. The Global Newton Method of GOVINDAN AND WILSON [6] is an example; those authors list the algorithms of LEMKE AND HOWSON [13], WILSON [21], and YAMAMOTO [22], among others, as special cases of their algorithm. These algorithms generally operate by perturbing or restricting the game in such a way that the modified game may be easily solved. Govindan and Wilson modify the payoffs sufficiently that the perturbed game has a unique equilibrium, then trace back to the original game using implications of the structure theorem of KOHLBERG AND MERTENS [12]. Yamamoto computes a proper equilibrium, as defined in MYERSON [17], of a normal form game by tracing out a path of profiles which change as the set of permitted mixed strategies is changed from the centroid to the

1. The Tracing Procedure will be capitalized to distinguish references to it from generic references to tracing a branch of a correspondence.

whole set.

GOVINDAN AND WILSON [7] note that approximation via quantal response equilibria differs from these homotopy approaches for computing equilibria. This paper fulfills the suggestion of McKelvey and Palfrey by presenting a homotopy approach to tracing branches of the logit equilibrium correspondence. This method of computing a Nash equilibrium has attractive properties for both theoretical and empirical applications. The trajectory of a branch of the logit equilibrium correspondence is governed by a dynamic process in λ which is a homotopy between the replicator dynamics and best-response dynamics. This interpretation parallels the description of the Tracing Procedure as a homotopy between an initial prior and final, consistent beliefs. The outputs of intermediate steps of an algorithm to trace the branch are logit equilibria useful in analysis of subject behavior in laboratory games; this algorithm, then, can be used as a component in programs to efficiently estimate values of λ from experimental data. The asymptotic behavior of the logit equilibrium correspondence yields estimates of the speed with which a path-following procedure will be able to give a good approximation to a Nash equilibrium. Finally, the homotopy characterization gives a tool for understanding some of the selection properties of quantal response equilibria in finite games.

This paper proceeds as follows. Section 2 introduces notation and summarizes the relevant properties of the logit equilibrium correspondence for normal form games. Section 3 derives the matrices used by the homotopy in following a branch of the correspondence. Section 4 gives a dynamic interpretation of a branch of the correspondence, and estimates the speed with which tracing a branch will provide a good approximation to a Nash equilibrium. Section 5 reports performance and timing results from an implementation of the method in Gambit. Section 6 applies the homotopy characterization to show that the principal branch of any quantal response equilibrium correspondence selects the risk-dominant equilibrium in 2×2 games if a natural monotonicity condition is satisfied. Section 7 summarizes, with suggestions for further development.

2 Quantal response equilibria in normal form games

The notation follows that of MCKELVEY AND PALFREY [15]. Given a finite n -player game in normal form with the set of players $N = \{1, \dots, n\}$, let S_i be the set of strategies available to each player $i \in N$, and let J_i be the number of strategies in S_i . A typical element of S_i is written s_{ij} , for $1 \leq j \leq J_i$. The set of strategy profiles S is the Cartesian product of the S_i , $S = \times_{i=1}^n S_i$. Let J_i be the number of strategies available to player i , and define $J = \sum_{i=1}^n J_i$. The payoff function for player i is $u_i: S \rightarrow \mathbb{R}$.

A mixed strategy for player i is a probability distribution over the set of i 's strategies; the set of such distributions for player i is denoted Δ_i . The set of mixed strategy profiles is $\Delta = \times_{i=1}^n \Delta_i$. A typical mixed strategy profile will be denoted by $\pi \in \Delta$, and the probability assigned to a strategy s_{ij} of player i is π_{ij} . The payoff functions for the players are extended over the set of mixed strategies in the usual expected value way.

Let (s_{ij}, π_{-i}) denote the mixed strategy profile where player i plays strategy s_{ij} with probability one, while all other players play according to the mixed strategy π . Then, define $\bar{u}: \Delta \rightarrow \times_{i=1}^n \mathbb{R}^{J_i}$ by

$$\bar{u}_{ij}(\pi) = u_i(s_{ij}, \pi_{-i}).$$

Here, $\bar{u}_{ij}(\pi)$ is the expected payoff to player i from playing his j th strategy, holding fixed his opponents' mixed strategies π_{-i} .

In the quantal response framework, players observe a noisy evaluation of the strategy values $\bar{u}_{ij}(\pi)$ of the form

$$\hat{u}_{ij}(\pi) = \bar{u}_{ij}(\pi) + \varepsilon_{ij},$$

where the ε_{ij} are random variables drawn according to some joint distribution. When the ε_{ij} are chosen independently from an extreme value distribution with parameter λ , a

logistic quantal response equilibrium profile (called a logit equilibrium below) is given by

$$\pi_{ij} = \frac{e^{\lambda \bar{u}_{ij}(\pi)}}{\sum_{k=1}^{J_i} e^{\lambda \bar{u}_{ik}(\pi)}}. \quad (1)$$

Therefore, the set of logit equilibria can be viewed as a correspondence mapping λ into a set of mixed strategy profiles in Δ . McKelvey and Palfrey show the correspondence is upper hemicontinuous, and has an odd number of members for almost all games and almost all λ , and its limit points as $\lambda \rightarrow \infty$ are Nash equilibria. Furthermore, it generically has a unique branch connecting the centroid of Δ to a unique Nash equilibrium, and generically connects up other limit Nash equilibrium points of the correspondence pairwise.

3 The homotopy approach

The computational problem is to trace out a branch of the logit equilibrium correspondence as a path of mixed strategy profiles π along with corresponding values of the parameter λ . To do this, the logit equilibria are expressed as the zeroes of a system of equations $H(\pi, \lambda) = 0$. To account for the possibility that the branch is not monotonic in λ , the path of (π, λ) pairs to be traced are parameterized by s ; that is, the homotopy will compute a parametric path $(\pi(s), \lambda(s))$, where s is interpreted as the arclength along the path.

Following GARCIA AND ZANGWILL [3], as presented in JUDD [11], the pair $y(s) = (\pi(s), \lambda(s))$, which has $J + 1$ components in total, satisfies the system of differential equations

$$\frac{dy_d}{ds} = (-1)^d C(\pi, \lambda) \left| \frac{\partial H}{\partial y}(\pi, \lambda) \right|_{-d} \quad \forall d = 1, \dots, J + 1, \quad (2)$$

where the vertical bars denote the determinant of a matrix, the notation A_{-d} means the matrix A with the d th column removed, and d indexes the components of $y(s)$.² The function $C(\pi, \lambda)$ is undetermined at this point, but is common across all components d .³ In other words, in tracing the parametric path, all that matters are the *ratios* $\frac{dy_d}{ds}/\frac{dy_{d'}}{ds}$ for $d \neq d'$. In keeping with the arclength interpretation of s , $C(\pi, \lambda)$ is chosen such that the side condition

$$\left[\sum_{i=1}^{J+1} \left(\frac{dy_d}{ds} \right)^2 \right]^{1/2} = 1$$

is satisfied. Given this side condition, the sign of C is not determined. This sign determines the orientation along the branch of the correspondence. When the starting point is $\lambda = 0$ with π equal to the centroid, $\pi_{ij} = \frac{1}{J_i}$, the sign is chosen to ensure $\frac{d\lambda}{ds} > 0$. It is straightforward to verify that the right-hand side of (2) is nonzero at this point.

The J equations comprising the system $H(\pi, \lambda)$ are indexed $H_{ij}(\pi, \lambda)$, corresponding to the j th strategy of player i . $H_{ij}(\pi, \lambda)$ is obtained by rearranging the logit equilibrium condition (1):

$$H_{ij}(\pi, \lambda) = e^{\lambda \bar{u}_{ij}(\pi)} - \pi_{ij} \sum_{k=1}^{J_i} e^{\lambda \bar{u}_{ik}(\pi)} = 0. \quad (3)$$

The partial derivatives of the equation $H_{ij}(\pi, \lambda)$ corresponding to strategy j of player i are divided into four cases.

2. This form of the solution path will be useful in the analysis which follows; practical details on the numerical linear algebra involved in efficiently following the path can be found in ALLGOWER AND GEORG [1].

3. The probabilities π_{ij} are assumed to be assigned to the components of y in increasing lexicographic order; that is, player 1's strategies are listed first, in ascending numerical order, followed by player 2's, and so on. The last $((J+1)$ st) component of the vector corresponds to λ .

Case 1 The derivative with respect to the corresponding probability, π_{ij} ,

$$\frac{\partial H_{ij}}{\partial \pi_{ij}} = - \frac{e^{\lambda \bar{u}_{ij}(\pi)}}{\pi_{ij}}. \quad (4)$$

Case 2 A derivative with respect to the probability π_{ik} of a strategy $k \neq j$ of player i ,

$$\frac{\partial H_{ij}}{\partial \pi_{ik}} = 0. \quad (5)$$

Case 3 A derivative with respect to the probability π_{lm} of a strategy m of player $l \neq i$,

$$\frac{\partial H_{ij}}{\partial \pi_{lm}} = e^{\lambda \bar{u}_{ij}(\pi)} \lambda \sum_{k=1}^{J_i} \left[\frac{\partial \bar{u}_{ij}(\pi)}{\partial \pi_{lm}} - \frac{\partial \bar{u}_{ik}(\pi)}{\partial \pi_{lm}} \right] \pi_{ik} \quad (6)$$

Case 4 The derivative with respect to λ ,

$$\frac{\partial H_{ij}}{\partial \lambda} = e^{\lambda \bar{u}_{ij}(\pi)} \sum_{k=1}^{J_i} [\bar{u}_{ij}(\pi) - \bar{u}_{ik}(\pi)] \pi_{ik} \quad (7)$$

In each case, since the homotopy is tracing points where $H(\pi, \lambda) = 0$, the definition of logit equilibrium (1) has been applied in expressing (4) through (7) in the given forms.

Observe that the differential equations (2) describing the branch of the logit equilibrium correspondence involves the matrix ∇H only via determinants where individual columns are removed. Therefore, if a row $(\nabla H)_{ij}$ is multiplied by a factor, all these determinants are multiplied by the same factor. So, defining a new matrix G with rows

$$G_{ij}(\pi, \lambda) = e^{-\lambda \bar{u}_{ij}(\pi)} \pi_{ij} (\nabla H(\pi, \lambda))_{ij},$$

the homotopy equations (2) can be expressed as

$$\frac{dy_d}{ds} = (-1)^d C(\pi, \lambda) \exp \left[\lambda \sum_{l=1}^n \sum_{j=1}^{J_l} \bar{u}_{lj}(\pi) \right] \left[\prod_{l=1}^n \prod_{j=1}^{J_l} \pi_{lj} \right]^{-1} |G(y)_{-d}|. \quad (8)$$

Since the function $C(\pi, \lambda)$ is arbitrary, the new factors can be folded into a redefined coefficient,

$$C'(\pi, \lambda) \equiv C(\pi, \lambda) \exp \left[\lambda \sum_{l=1}^n \sum_{j=1}^{J_l} \bar{u}_{lj}(\pi) \right] \left[\prod_{l=1}^n \prod_{j=1}^{J_l} \pi_{lj} \right]^{-1},$$

which produces a final version of the homotopy equation

$$\frac{dy_d}{ds} = (-1)^d C'(\pi, \lambda) |G(\pi, \lambda)_{-d}| \forall d = 1, \dots, J+1, \quad (9)$$

To summarize, defining a $J \times (J+1)$ matrix G , with rows indexed by the set of strategies $\bigcup_{i=1}^n S_i$ and columns indexed by $[\bigcup_{i=1}^n S_i] \cup \{\lambda\}$, the set of points $y(s) = (\pi(s), \lambda(s))$ on a branch of the logit equilibrium correspondence is characterized by the system of differential equations (9), where

$$G_{\pi_{ij}, \pi_{ij}} = 1 \forall i \in N, j = 1, \dots, J_i \quad (10)$$

$$G_{\pi_{ij}, \pi_{ik}} = 0 \forall i \in N, k \neq j \quad (11)$$

$$G_{\pi_{ij}, \pi_{lm}} = \lambda \pi_{ij} \sum_{k=1}^{J_i} \left[\frac{\partial \bar{u}_{ij}(\pi)}{\partial \pi_{lm}} - \frac{\partial \bar{u}_{ik}(\pi)}{\partial \pi_{lm}} \right] \pi_{ik} \\ \forall i \in N, j = 1, \dots, J_i, l \neq i, m = 1, \dots, J_l \quad (12)$$

$$G_{\pi_{ij}, \lambda} = \pi_{ij} \sum_{k=1}^{J_i} [\bar{u}_{ij}(\pi) - \bar{u}_{ik}(\pi)] \pi_{ik} \forall i \in N, j = 1, \dots, J_i \quad (13)$$

4 Qualitative characteristics of algorithm behavior

4.1 A dynamic interpretation of the principal branch

HARSANYI AND SELTEN [8] describe their Tracing Procedure as a model of expectation formation and strategy choice where players form and adapt beliefs and tentative choices

over an interval of time. The correspondence consisting of paths followed by the Tracing Procedure has a similar structure to the logit equilibrium correspondence. As an additional parallel, the branches of the logit equilibrium correspondence can be described in dynamic terms.

Suppose the parameter λ is interpreted as time, and consider the application of logit equilibrium to a decision problem. Then, the logit equilibrium correspondence can be interpreted in terms of the replicator dynamics (TAYLOR AND JONKER [19]). The replicator dynamics evolves the proportion of a population playing a strategy s_{ij} over time according to the equation

$$\frac{d\pi_{ij}(t)}{dt} = \pi_{ij}(t) \left[\bar{u}_{ij}(\pi(t)) - \sum_{k=1}^{J_i} \pi_{ik}(t) \bar{u}_{ik}(\pi(t)) \right]$$

Theorem 1. *In a decision problem, the logit equilibrium correspondence consists of a single branch, which coincides with the the path taken by the replicator dynamics when started at the centroid.*

Proof. Begin by noting that

$$\begin{aligned} \frac{d\pi_{ij}}{d\lambda} &= \frac{d\pi_{ij}(s)}{ds} \left(\frac{d\lambda(s)}{ds} \right)^{-1} \\ &= \frac{C'(\pi, \lambda)|_{G_{-\pi_{ij}}}}{C'(\pi, \lambda)|_{G_{-\lambda}}} \\ &= \frac{\pi_{ij} \sum_{k=1}^{J_i} [\bar{u}_{ij}(\pi) - \bar{u}_{ik}(\pi)] \pi_{ik}}{1}, \end{aligned}$$

with the last equality following because there are no cross-player terms $G_{\pi_{ij}, \pi_{lm}}$ in a decision problem. Therefore, the logit equilibrium path is the same as the path followed by replicator dynamics when started at the centroid, with λ playing the role of time. Unique-

ness follows since two strategies must be played with equal probability if they have the same payoff. \square

Now consider the case of a proper game with two or more players. Writing the ratio of the probabilities that two strategies s_{ij} and s_{ik} are played in a logit equilibrium, it follows that⁴

$$\frac{\pi_{ij}}{\pi_{ik}} = \exp [\lambda(\bar{u}_{ij} - \bar{u}_{ik})] \quad (14)$$

$$\begin{aligned} \frac{d}{d\lambda} \left[\frac{\pi_{ij}}{\pi_{ik}} \right] &= \exp [\lambda(\bar{u}_{ij} - \bar{u}_{ik})] \times \left\{ \bar{u}_{ij} - \bar{u}_{ik} + \lambda \sum_{l \neq i} \sum_{m=1}^{J_l} \left(\frac{\partial \bar{u}_{ij}}{\partial \pi_{lm}} - \frac{\partial \bar{u}_{ik}}{\partial \pi_{lm}} \right) \frac{\partial \pi_{lm}}{\partial \lambda} \right\} \\ &= \frac{\pi_{ij}}{\pi_{ik}} \times \left\{ \bar{u}_{ij} - \bar{u}_{ik} + \lambda \sum_{l \neq i} \sum_{m=1}^{J_l} \left(\frac{\partial \bar{u}_{ij}}{\partial \pi_{lm}} - \frac{\partial \bar{u}_{ik}}{\partial \pi_{lm}} \right) \frac{\partial \pi_{lm}}{\partial \lambda} \right\} \end{aligned} \quad (15)$$

Again interpreting λ in the role of time, equation (15) expresses the dynamics implied by increasing λ as changing smoothly from the replicator dynamics to best-reply dynamics. At $\lambda = 0$, equation (15) reduces to the replicator dynamics. The expression in square brackets captures the change in the payoff difference between strategies s_{ij} and s_{ik} , as players other than i change their strategies. Increasing λ , then, increases the speed with which player i reacts to changes in opponents' play. As λ goes to infinity, this amounts to best-response dynamics.

The analog of the principal branch of the logit equilibrium correspondence in the Tracing Procedure is the distinguished path. Following the distinguished path is interpreted as tracing out a chain of introspection among rational players facing strategic uncertainty, which chain leads in the end to a Nash equilibrium. This corresponds to the blending of what Harsanyi and Selten call first-order and second-order information avail-

4. For brevity, $\pi_{ij}(\lambda)$ will be abbreviated π_{ij} , and $\bar{u}_{ij}(\pi(\lambda))$ will be abbreviated \bar{u}_{ij} . That is to say, unless otherwise noted, these are evaluated along a branch of the correspondence.

able to the players. First-order information expresses players' information about each others' likely choices, whereas second-order information pertains to each others' reactions to first-order information. At a Nash equilibrium, these must be consistent, and the Tracing Procedure accomplishes this by gradually feeding second-order information into the calculation.

An analogous structure can be seen in equation (15). Players' first-order information is captured by the difference in payoffs $\bar{u}_{ij}(\pi) - \bar{u}_{ik}(\pi)$, while second-order information is expressed in the term within square brackets. As the path is traversed, players initially update their play naively based upon their first-order information according to the replicator dynamics; gradually, second-order information is fed into the adjustment process through the increase in λ .

This interpretation suffers from weaknesses similar to those in the interpretation of the Tracing Procedure. Even in generic games, the principal branch may have turning points, leading to intervals on which λ is decreasing while following the principal branch in the direction from the centroid at $\lambda = 0$ to the limiting Nash equilibrium. A similar phenomenon arises on the distinguished path of the Tracing Procedure, where segments moving “backwards in time” exist even in generic games. Nongenerically, branch points may occur in the logit equilibrium correspondence, just as they may in the Tracing Procedure. Harsanyi and Selten introduce the logarithmic version of the Tracing Procedure to remove this degeneracy. For the logit equilibrium correspondence, standard methods in numerical path following are available to characterize and address these cases.

4.2 Asymptotic behavior and efficient tracing

Application of the logit equilibrium approach to the problem of computing a Nash equilibrium requires analysis of the rate at which logit equilibria converge to a Nash equilibrium. This section outlines properties of this convergence as λ grows large. The key results are

that the change in the logit equilibrium profile slows faster than λ^{-1} , and that an algorithm to trace the path can take exponentially increasing steps in λ while tracing the path and still maintain an assurance of approximating the path with security.

Lemma 2. *There exists some $\lambda^* < \infty$ such that there are no turning points in the logit equilibrium correspondence for $\lambda > \lambda^*$.*

Proof. At a turning point in a branch of the correspondence, $|G_{-\lambda}| = 0$. This determinant is a nontrivial polynomial in $J + 1$ variables: λ and the J probabilities in π . So, there are at most a finite number of candidate points $(\pi, \lambda) \in \Delta \times [0, \infty)$ which could be turning points (and will be if they lie on one of the branches of the correspondence). Since the number of such points is finite, there must be one with the largest λ ; denote that value λ^* . □

In view of Lemma 2, all asymptotic results in this section consider only the region of the branch being traced where $\lambda > \lambda^*$.

Lemma 3. *Suppose that π^* is a limiting logit equilibrium such that $\bar{u}_{ik}(\pi^*) < \bar{u}_{ij}(\pi^*)$ for some strategies s_{ij} and s_{ik} of player i . Then, $\lim_{\lambda \rightarrow \infty} \lambda \pi_{ik}(\lambda) = 0$.*

Proof. Without loss of generality, suppose that $\pi_{ij}^* > 0$. Then, by manipulation of the identity (14),

$$\begin{aligned} \pi_{ij} &= \pi_{ik} e^{\lambda[\bar{u}_{ij} - \bar{u}_{ik}]} \\ \pi_{ij}^{\pi_{ik}} &= \pi_{ik}^{\pi_{ik}} e^{\lambda \pi_{ik} [\bar{u}_{ij} - \bar{u}_{ik}]} \\ \pi_{ik} \ln \pi_{ij} &= \pi_{ik} \ln \pi_{ik} + \lambda \pi_{ik} [\bar{u}_{ij} - \bar{u}_{ik}] \end{aligned}$$

Since $\pi_{ik} \rightarrow 0$ and π_{ij} converges to a positive limit, the left-hand side goes to zero as $\lambda \rightarrow \infty$. Also since $\pi_{ik} \rightarrow 0$, $\pi_{ik} \ln \pi_{ik} \rightarrow 0$ as $\lambda \rightarrow \infty$. Therefore, $\lambda \pi_{ik} [\bar{u}_{ij} - \bar{u}_{ik}] \rightarrow 0$ as $\lambda \rightarrow \infty$.

Since it is assumed that $\bar{u}_{ij}(\pi^*) - \bar{u}_{ik}(\pi^*) > 0$, it must be that $\lambda\pi_{ik} \rightarrow 0$ as $\lambda \rightarrow \infty$, completing the proof. \square

One way to express how close a profile π is to being a Nash equilibrium is to evaluate the velocity of the replicator dynamics evaluated at π . Lemma 4 shows that, along a branch of the logit equilibrium correspondence, this velocity goes to zero linearly in λ^{-1} .

Lemma 4. *Along any branch of the logit equilibrium correspondence,*

$$\lim_{\lambda \rightarrow \infty} \left| \lambda \pi_{ij} \sum_{k=1}^{J_i} [\bar{u}_{ij} - \bar{u}_{ik}] \pi_{ik} \right| < \infty.$$

Proof. Taking the logarithm of the identity (14) produces

$$\ln \left(\frac{\pi_{ij}}{\pi_{ik}} \right) = \lambda [\bar{u}_{ij} - \bar{u}_{ik}].$$

Therefore,

$$\begin{aligned} \lambda \pi_{ij} \sum_{k=1}^{J_i} [\bar{u}_{ij} - \bar{u}_{ik}] \pi_{ik} &= \pi_{ij} \sum_{k=1}^{J_i} \lambda [\bar{u}_{ij} - \bar{u}_{ik}] \pi_{ik} \\ &= \pi_{ij} \sum_{k=1}^{J_i} \pi_{ik} \ln \left(\frac{\pi_{ij}}{\pi_{ik}} \right) \\ &= \pi_{ij} \left[\ln \pi_{ij} - \sum_{k=1}^{J_i} \pi_{ik} \ln \pi_{ik} \right]. \end{aligned}$$

As $\lambda \rightarrow \infty$, the right-hand side converges to a finite limit. \square

The next theorem provides the key result for the convergence behavior of logit equilibria.

Theorem 5. *Along a branch of the logit equilibrium correspondence, $\lim_{\lambda \rightarrow \infty} \lambda \frac{d\pi_{ij}}{d\lambda} = 0$.*

Proof. Begin by noting that

$$\lambda \frac{d\pi_{ij}}{d\lambda} = \lambda \frac{|G_{-\pi_{ij}}|}{|G_{-\lambda}|}. \quad (16)$$

The matrices $G_{-\pi_{ij}}$ and $G_{-\lambda}$ differ in only one column. The matrix $G_{-\pi_{ij}}$ contains the column $G_{\pi_{ij}, \lambda}$, whereas the matrix $G_{-\lambda}$ contains a column with elements of the forms $G_{\pi_{ij}, \pi_{ij}}$, $G_{\pi_{ij}, \pi_{ik}}$, and $G_{\pi_{ij}, \pi_{lm}}$. Without loss of generality, the differing column can be thought to have the same column index, since permuting the columns of a matrix at most changes the sign of the determinant.

To establish the result, it will be argued that the numerator can be written in the form

$$\lambda |G_{-\pi_{ij}}| = \sum_{d=0}^{J-1} a_d(\pi, \lambda) \lambda^d$$

and the denominator in the form

$$|G_{-\lambda}| = \sum_{d=0}^J b_d(\pi) \lambda^d,$$

where the functions $a_d(\pi, \lambda)$ and $b_d(\pi)$ are bounded.

Recall that the determinant of a $J \times J$ matrix can be expressed as a sum of $J!$ terms, where each term is the product of J elements of the matrix, chosen such that each row and each column of the matrix is represented exactly once. Since λ only appears in the matrix G through multiplication of matrix elements of the form $G_{\pi_{ij}, \pi_{lm}}$, this sum can be organized according to the number of those elements appearing in each term. In $G_{-\lambda}$ there exist products of up to J elements of the form $G_{\pi_{ij}, \pi_{lm}}$. In $G_{-\pi_{ij}}$ there exist only products of up to $J - 1$ elements of that form, as the column corresponding to π_{ij} is

removed and replaced with a column of elements of the form $G_{\pi_{ij},\lambda}$.

Collecting terms in this way, denote the coefficient of λ^d in the sum corresponding to $|G_{-\pi_{ij}}|$ by $\tilde{a}_d(\pi)$ for $d = \{0, 1, \dots, J-1\}$, and denote the coefficient of λ^d in the sum corresponding to $|G_{-\lambda}|$ by $b_d(\pi)$ for $d = \{0, 1, \dots, J\}$. With this notation, the numerator can be expressed as $\lambda|G_{-\pi_{ij}}| = \sum_{d=1}^J \tilde{a}_{d-1}(\pi)\lambda^d$, and the denominator as $|G_{-\lambda}| = \sum_{d=0}^J b_d(\pi)\lambda^d$. Next, define $a_d(\pi, \lambda) = \lambda\tilde{a}_d(\pi)$. With this definition, the numerator can be reexpressed in the form $\lambda|G_{-\pi_{ij}}| = \sum_{d=0}^{J-1} a_d(\pi, \lambda)\lambda^{d-1}$.

The next step is to establish that the $a_d(\pi, \lambda)$ and $b_d(\pi)$ functions so defined are bounded. The $b_d(\pi)$ functions consist of sums of a finite number of terms which are products of matrix elements of the forms $G_{\pi_{ij},\pi_{lm}}$ and $G_{\pi_{ij},\lambda}$. Since the presence of λ is removed by the way terms were collected in defining $b_d(\pi)$, these terms are clearly bounded, and so the sum is bounded. The argument for the boundedness of $a_d(\pi, \lambda)$ being bounded requires an extra step, as $\lambda \rightarrow \infty$. Note that every term in the sum inside $\tilde{a}_d(\pi)$ must by definition contain an expression of the $G_{\pi_{ij},\lambda}$ form as a multiplicand. Lemma 4 shows that the quantity $\lambda G_{\pi_{ij},\lambda}$ converges to a finite limit as $\lambda \rightarrow \infty$. The remainder of each term depends only on π , and so is bounded. Since $\lambda G_{\pi_{ij},\lambda}$ is continuous in both π and λ , each term in the sum is bounded for all π and λ , and so $a_d(\pi, \lambda)$ is bounded.

If the branch of the correspondence under consideration limits to a totally mixed equilibrium, the proof is complete, since (generically) $b_J(\pi) \not\rightarrow 0$ along such a branch, and the difference in order in λ dominates. However, if some strategies are strictly inferior in the limit, then $b_J(\pi) \rightarrow 0$ and the order difference argument requires an extra step. Pick some d , and note that for a term to appear in $b_d(\pi)$, it must contain exactly d of the strategy probabilities π_{ij} as products, coming from the multiplication of $G_{\pi_{ij},\pi_{lm}}$ elements. Recall that the matrix $G_{-\lambda}$ differs from the matrix $G_{-\pi_{ij}}$ in only one column. Therefore, for any given term in the sum inside $b_d(\pi)$, there is a corresponding term in the sum inside $\tilde{a}_{d-1}(\pi)$ which is obtained by multiplying the entries with the same row and column

indices. Therefore, the term inside $\tilde{a}_{d-1}(\pi)$ also has exactly those same d probabilities as multiplicands. So, if a product in the sum inside $b_d(\pi)$ is tending to zero in the limit because it contains a factor π_{ij} corresponding to a strategy not played in the limit, there is a corresponding product in the sum inside $a_{d-1}(\pi)$ also tending to zero because it contains that same factor. Because of this one-to-one relationship between the terms inside $b_d(\pi)$ and $a_{d-1}(\pi)$, if $b_d(\pi) \rightarrow 0$, then $a_{d-1}(\pi) \rightarrow 0$. If J^* is the number of strategies played with positive probability in the limiting equilibrium, then the leading coefficients not going to zero are $b_{J^*}(\pi)$ and $a_{J^*-1}(\pi, \lambda)$. The denominator is thus of higher order than the numerator, and the result follows. □

Theorem 5 now permits a characterization of the speed with which a path-following algorithm will traverse a branch of the logit equilibrium correspondence and converge to a Nash equilibrium. Typically, the path defined by (9) is traversed by a two-step procedure known as a predictor-corrector method. The predictor phase is a numerical integration step, advancing from a logit equilibrium with parameter λ to one with parameter $\lambda + h$ using a first-order expansion of the form

$$\pi_{ij}(\lambda + h) = \pi_{ij}(\lambda) + h \frac{d\pi_{ij}}{d\lambda}(\lambda). \quad (17)$$

While an implementation could simply treat (9) as a differential equation and use only numerical integration methods like (17) to trace the path, this would ignore the information that the differential equation (9) characterizes the solution of a set of equations. Therefore, a corrector step is then performed, which uses the contractive properties of Newton's method for finding a zero of a system of equations for refining the accuracy of the new point.

The choice of the steplength h is important in practical application. If h is small, many steps are required to traverse the path; if h is too large, there is a risk of exceeding the convergence radius of Newton's method and diverging from the path. Strategies exist

for choosing h adaptively, based upon the convergence rate of Newton's method in the corrector step, which in turn depends on the accuracy of initial guess from the predictor step. These strategies generally attempt to target a desired convergence rate, chosen as a parameter by the user. Adaptive steplength choice is of particular importance for tracing the logit equilibrium correspondence, as the parameter λ is not bounded and so traversing the path with a fixed choice of h would take prohibitively long to reach an acceptable approximation of the limiting Nash equilibrium. The following result suggests that, as λ grows large, the optimal steplength is roughly proportional to λ .

Theorem 6. *Fix a constant $c > 1$. Then, the distance between $\pi(c\lambda)$ and the value estimated using the predictor strategy (17) with step $h = c\lambda$ limits to a constant as $\lambda \rightarrow \infty$.*

Proof. Strategies which are strictly inferior in the limiting equilibrium are played with exponentially decreasing probability as $\lambda \rightarrow \infty$. So, without loss of generality, consider a branch of the correspondence which has a totally mixed Nash equilibrium as its limit. The argument is more conveniently expressed in terms of the ratios of probabilities. Equation (15) expresses how the ratio of the probabilities of two strategies for player i change as λ changes, given how strategies for other players change with λ . Taking $\pi_{-i}(\lambda)$ as given, this is a first-order linear differential equation for the ratio $\rho = \frac{\pi_{ij}}{\pi_{ik}}$, which has solution

$$\rho(\lambda) = \exp \left[\int_{\lambda^*}^{\lambda} [\bar{u}_{ij}(\pi(t)) - \bar{u}_{ik}(\pi(t))] dt + \int_{\lambda^*}^{\lambda} t f(t) dt \right],$$

where $f(\lambda) = \sum_{l \neq i} \sum_{m=1}^{J_l} \left(\frac{\partial \bar{u}_{ij}}{\partial \pi_{lm}} - \frac{\partial \bar{u}_{ik}}{\partial \pi_{lm}} \right) \frac{d\pi_{lm}}{d\lambda}$, and where the branch is considered only for values of λ greater than the λ^* defined in Lemma 2, so that $\rho(\lambda)$ is meaningful. Therefore,

$$\rho(c\lambda) = \rho(\lambda) \times \exp \left[\int_{\lambda}^{c\lambda} [\bar{u}_{ij}(\pi(t)) - \bar{u}_{ik}(\pi(t))] dt + \int_{\lambda}^{c\lambda} t f(t) dt \right].$$

Meanwhile, let the logit equilibrium estimated by predictor (17) at $c\lambda$ be denoted by

$$\tilde{\rho}(c\lambda) = \rho(\lambda) + (c-1)\lambda\rho'(\lambda).$$

Therefore, the error in the estimate is given by

$$\rho(c\lambda) - \tilde{\rho}(c\lambda) = \rho(\lambda) \left\{ \exp \left[\int_{\lambda}^{c\lambda} (\bar{u}_{ij}(\pi(t)) - \bar{u}_{ik}(\pi(t))) dt + \int_{\lambda}^{c\lambda} t f(t) dt \right] - 1 \right\} - \lambda\rho'(\lambda).$$

The first integral in the exponential tends to a constant as $\lambda \rightarrow \infty$, as

$$\begin{aligned} \frac{d}{d\lambda} \int_{\lambda}^{c\lambda} (\bar{u}_{ij}(\pi(t)) - \bar{u}_{ik}(\pi(t))) dt &= [\bar{u}_{ij}(\pi(c\lambda)) - \bar{u}_{ik}(\pi(c\lambda))] - [\bar{u}_{ij}(\pi(\lambda)) - \bar{u}_{ik}(\pi(\lambda))] \\ &= \frac{1}{c\lambda} \ln \rho(c\lambda) - \frac{1}{\lambda} \ln \rho(\lambda), \end{aligned}$$

which has a limit of zero. Similarly,

$$\frac{d}{d\lambda} \int_{\lambda}^{c\lambda} t f(t) dt = c\lambda f(c\lambda) - \lambda f(\lambda).$$

Theorem 5 implies that $\lambda f(\lambda) \rightarrow 0$ as $\lambda \rightarrow \infty$. Therefore, the difference $\rho(c\lambda) - \tilde{\rho}(c\lambda)$ tends to a constant as $\lambda \rightarrow \infty$. \square

This result says that if one were to use an exponentially increasing stepsize in the predictor step, the error in the initial guess in the predictor step given by (17) would not blow up. Therefore, traversing the path “faster” in λ as λ increases would not compromise the ability to follow the path securely. In fact, techniques that adaptively choose the stepsize based on the accuracy of the predictor step, such as the one used in the implementation described in the next section, generate such rapidly increasing stepsizes. While the implementation tests in the next section suggest the method yields a very good approxi-

mation to a Nash equilibrium quickly, Theorem 6 indicates that this is in fact a general theoretical property of the algorithm.

5 Implementation of the algorithm

Theorems 5 and 6 suggest that a method of tracing the logit equilibrium correspondence will have good asymptotic properties in computing an approximate Nash equilibrium. They leave unaddressed two issues, which could have a negative impact on the usefulness of this method for computing Nash equilibrium:

1. The results are only asymptotic. There is no characterization of the total arclength along the path that would need to be traversed in order to reach the region where the asymptotic results apply. Further, this arclength could increase with the size of the game J as, in particular, the potential total number of turning points in the correspondence increases with J .
2. The actual computational cost in processor time may scale poorly in J , making the method impractical even for moderate-sized games. In particular, tracing the path involves linear algebra on $J + 1$ matrices of size $J \times J$.

The method has been implemented in Gambit,⁵ beginning in version 0.97.0.1. The procedure was evaluated on a sample of games with different numbers of players and strategies. For each size of game, 1,000 games were generated at random, where each payoff entry is drawn iid from the uniform distribution on $[0, 1]$.

5. Gambit (<http://econweb.tamu.edu/gambit>) is an open source software project, founded by Richard McKelvey, for computation in finite games.

	$J_i = 2$	$J_i = 3$	$J_i = 4$	$J_i = 5$	$J_i = 10$	$J_i = 20$
$n = 2$	270 (158,369)	339 (158,589)	397 (158,680)	456 (253,708)	704 (449,1137)	1136 (749,3418)
$n = 3$	304 (158,513)	384 (158,673)	467 (158,762)	552 (310,1053)		
$n = 4$	302 (158,634)	429 (158,804)	545 (245,1216)			
$n = 5$	333 (158,642)	483 (158,991)				

Table 1. Number of steps required to reach $\lambda = 1,000,000$. In each cell, the number on top is the median number of steps over the sample of games; the range in parentheses reflects the minimum and maximum numbers.

	$\lambda = 10^1$	$\lambda = 10^2$	$\lambda = 10^3$	$\lambda = 10^4$	$\lambda = 10^5$
$n = 2, J_i = 2$	37 (37,39)	107 (60,157)	189 (85,413)	219 (103,451)	243 (133,476)
$n = 2, J_i = 3$	37 (37,49)	169 (60,337)	266 (85,541)	300 (109,582)	325 (133,607)
$n = 2, J_i = 4$	37 (37,55)	242 (60,423)	387 (173,761)	428 (208,953)	454 (232,981)
$n = 2, J_i = 20$	37 (37,41)	504 (140,995)	985 (674,1591)	1077 (700,1716)	1111 (724,1743)
$n = 3, J_i = 2$	37 (37,76)	118 (60,252)	221 (85,435)	255 (109,462)	280 (133,487)
$n = 4, J_i = 2$	37 (37,62)	121 (60,304)	218 (85,534)	254 (109,655)	279 (133,683)
$n = 5, J_i = 2$	37 (37,70)	125 (60,399)	250 (85,586)	285 (109,630)	310 (133,661)

Table 2. Number of steps required to reach $\lambda = 10^k$. In each cell, the number on top is the median number of steps over the sample of games; the range in parentheses reflects the minimum and maximum numbers.

Table 1 presents statistics on the number of steps of the homotopy procedure required to reach $\lambda = 1,000,000$ on the principal branch. The number of steps required increases modestly in the size of the game; typically, the principal branch does not become more

complicated, in the sense of having more turning points, as the game size increases. The worst-case step counts correspond, in general, to games where the principal branch does have a turning point. The observation that the minimum number of steps is equal to 158 for several of the cells results from an implementation restriction on the rate at which the steplength may be increased; for games requiring 158 steps from $\lambda = 0$ to $\lambda = 1,000,000$, the implementation increases the steplength by the maximum permitted each step.

Theorems 5 and 6 imply that as λ grows, the algorithm should take larger and larger steps in the λ dimension as the branch asymptotes quickly towards the limiting equilibrium profile. Table 2 breaks down the progress of the procedure towards larger values of λ for selected classes of games. After approximately $\lambda = 10^3$, progress appears exponential; a given number of steps advances λ by an order of magnitude. For example, with $n = 2$ and $J_i = 2$, the median number of steps to $\lambda = 10^4$ is 30 greater than the median to $\lambda = 10^3$; the median for $\lambda = 10^5$ is then only 24 steps greater than for $\lambda = 10^4$. The worst-case numbers of steps show a similar pattern; these indicate that most of the segments containing turning points and other features which impede computation are contained within the $\lambda \in [10, 100]$ interval.

While the total number of steps in the tracing does not increase dramatically, the computational cost of each step will increase. Total processor time, then, will be increasing, even when the number of homotopy steps does not. Table 3 summarizes total processing time, which accounts for all operations required to trace the branch.⁶ Within the classes of games under consideration, the running time scales well on typical games. Once again, the worst-case times increase more rapidly; these again correspond to games where the principal branch either has or almost has a turning point, requiring a shorter step size and more steps to follow the branch securely.

6. These timings were generated on a Linux workstation with 1.7 GHz Xeon processors. The implementation in Gambit version 0.97.0.4 was used, compiled with gcc version 3.2.3, with optimization (-O2) and debugging symbols (-g).

	$J_i = 2$	$J_i = 3$	$J_i = 4$	$J_i = 5$	$J_i = 10$	$J_i = 20$
$n = 2$	0.05 (0.03,0.08)	0.08 (0.05,0.19)	0.13 (0.08,0.26)	0.19 (0.12,0.39)	1.03 (0.66,3.18)	12.85 (7.78,45.80)
$n = 3$	0.08 (0.05,0.19)	0.20 (0.14,0.79)	0.47 (0.30,2.11)	1.03 (0.62,4.10)		
$n = 4$	0.15 (0.10,0.56)	0.68 (0.40,3.58)	2.63 (1.56,15.70)			
$n = 5$	0.32 (0.22,1.60)	2.93 (1.81,20.50)				

Table 3. Time (in seconds) required to reach $\lambda = 1, 000, 000$. In each cell, the number on top is the median time over the sample of games; the range in parentheses reflects the minimum and maximum times.

6 Application: Selection of risk-dominant equilibria

Since the logit equilibrium correspondence generically has a branch connecting the centroid at $\lambda = 0$ to a unique Nash equilibrium, McKelvey and Palfrey suggest using this limiting equilibrium as a means of selecting from the set of Nash equilibria. Harsanyi and Selten suggest that a desirable property for a solution concept is selection of a risk-dominant equilibrium when one exists. In a 2×2 game with two strict equilibria, risk dominance implies that the strategy played by a player in the risk-dominant equilibrium is a best reply against the mixed strategy where his opponent plays both strategies with equal probability. The next theorem shows that, if an intuitive monotonicity assumption holds, the principal branch of *any* quantal response equilibrium correspondence selects the risk dominant equilibrium in this class of games.

Let $f_j(\bar{u}_i, \lambda)$ be any quantal response function, assigning a probability to player i 's strategy s_{ij} as a function of the vector of payoffs $\bar{u}_i = (\bar{u}_{ij})_{i=1}^{J_i}$ to all strategies of player i , and the parameter λ . Assume that λ is chosen such that increasing λ corresponds to increasing precision, or decreasing noise, and that λ ranges from $\underline{\lambda}$ to $\bar{\lambda}$, with $\underline{\lambda}$ corresponding to “infinite noise.”

Theorem 7. *Suppose a quantal response function f_j satisfies the following:*

1. *The derivatives $\frac{\partial f_j}{\partial u_k}$ and $\frac{\partial f_j}{\partial \lambda}$ are all continuous;*
2. *$\frac{\partial f_j}{\partial u_j} \geq 0$ (monotonicity in payoff of the strategy), with equality only when $\lambda = \underline{\lambda}$ (infinite noise);*
3. *$\frac{\partial f_j}{\partial \lambda} > \frac{\partial f_k}{\partial \lambda}$ if $\bar{u}_{ij} > \bar{u}_{ik}$ (decreasing noise);*
4. *The limit points of branches of the correspondence as $\lambda \rightarrow \bar{\lambda}$ are Nash equilibria.*

Then, the principal branch of the quantal response correspondence limits to the risk-dominant equilibrium in 2×2 games with two strict Nash equilibria.

Proof. Suppose that the strict Nash equilibria are at (s_{11}, s_{21}) and (s_{12}, s_{22}) , with the former being risk-dominant, and let a_{jk}^i denote the payoff to player i when player 1 plays s_{1j} and player 2 plays s_{2k} . The quantal response probabilities are given by $\pi_{ij} = f_j(\bar{u}_{11}, \bar{u}_{12}, \lambda)$. Since the probabilities for a player must sum to one, the correspondence can be expressed as the solution to the pair of equations

$$H_1 \equiv \pi_{11} - f_1(\bar{u}_{11}, \bar{u}_{12}, \lambda) = 0$$

$$H_2 \equiv \pi_{21} - f_2(\bar{u}_{21}, \bar{u}_{22}, \lambda) = 0$$

The partial derivatives of interest for H_1 are

$$\frac{\partial H_1}{\partial \pi_{11}} = 1$$

$$\begin{aligned} \frac{\partial H_1}{\partial \pi_{21}} &= - \left\{ \frac{\partial f_1}{\partial u_{11}} (a_{11}^1 - a_{12}^1) + \frac{\partial f_1}{\partial u_{12}} (a_{21}^1 - a_{22}^1) \right\} \\ &= \frac{\partial f_1}{\partial u_{11}} (a_{11}^1 - a_{21}^1 + a_{22}^1 - a_{12}^1) \end{aligned}$$

$$\frac{\partial H_1}{\partial \lambda} = - \frac{\partial f_1}{\partial \lambda}.$$

Because (s_{11}, s_{21}) and (s_{12}, s_{22}) are strict Nash equilibria, the quantity

$$C_1 \equiv (a_{11}^1 - a_{21}^1 + a_{22}^1 - a_{12}^1)$$

is positive. The expressions for the partial derivatives of H_2 are analogous.

Now, to implement the homotopy in (2), compute

$$\begin{aligned} \frac{d\pi_{11}}{ds} &= (-1) \begin{vmatrix} \frac{\partial H_1}{\partial \pi_{21}} & \frac{\partial H_1}{\partial \lambda} \\ \frac{\partial H_2}{\partial \pi_{21}} & \frac{\partial H_2}{\partial \lambda} \end{vmatrix} \\ &= -C_1 \frac{\partial f_1}{\partial u_{11}}(\bar{u}_{11}, \bar{u}_{12}, \lambda) \frac{\partial f_1}{\partial \lambda}(\bar{u}_{21}, \bar{u}_{22}, \lambda) - \frac{\partial f_1}{\partial \lambda}(\bar{u}_{11}, \bar{u}_{12}, \lambda). \end{aligned} \quad (18)$$

At $s = 0$, this right hand side of (18) is strictly negative. Because (s_{11}, s_{21}) is risk-dominant, it must be that $\bar{u}_{11} > \bar{u}_{12}$ when player 2 randomizes with equal probability between his strategies. Therefore, by assumption on f_j , $\frac{\partial f_1}{\partial \lambda} > 0$. A similar calculation establishes that $\frac{d\pi_{21}}{ds} < 0$ at $s = 0$.

Next, consider the change in λ along the path. This is calculated by

$$\begin{aligned} \frac{d\lambda}{ds} &= (-1) \begin{vmatrix} \frac{\partial H_1}{\partial \pi_{11}} & \frac{\partial H_1}{\partial \pi_{21}} \\ \frac{\partial H_2}{\partial \pi_{11}} & \frac{\partial H_2}{\partial \pi_{21}} \end{vmatrix} \\ &= -1 + \left[C_1 \frac{\partial f_1}{\partial u_{11}}(\bar{u}_{11}, \bar{u}_{12}, \lambda) \right] \left[C_2 \frac{\partial f_1}{\partial u_{21}}(\bar{u}_{21}, \bar{u}_{22}, \lambda) \right], \end{aligned}$$

where $C_2 > 0$ is the analogue from player 2's perspective of C_1 . At $s = 0$, this evaluates to -1 .

Since all three of $\frac{d\pi_{11}}{ds}$, $\frac{d\pi_{12}}{ds}$, and $\frac{d\lambda}{ds}$ are negative, to traverse the branch from $\lambda = \underline{\lambda}$, choose the negative orientation of the curve; therefore, at $\lambda = \underline{\lambda}$, both π_{11} and π_{21} are strictly increasing in increasing λ , and so there exists some $\sigma > 0$ such that $\pi_{11}(\sigma) > \frac{1}{2}$, $\pi_{21}(\sigma) > \frac{1}{2}$, and $\lambda(\sigma) > 0$ after traversing an arclength of σ on the branch. At such a point, the expression on the right-hand side of (18) remains negative. Since the branch is being traversed in the negative orientation, this corresponds to a positive change in π_{11} as the branch is traversed in the direction of increasing λ ; the argument for π_{21} continuing to

increase is parallel. Therefore, π_{11} and π_{21} are increasing as the principal branch is traversed in the direction away from the centroid at $\lambda = \underline{\lambda}$. Since the expression in (18) is strictly negative, the path cannot pass through a branch point in the correspondence, since a branch point implies all the derivatives in (2) vanish.

Finally, the branch must converge to a Nash equilibrium as $\lambda \rightarrow \bar{\lambda}$. Since (s_{11}, s_{21}) is risk-dominant, the mixed-strategy equilibrium in the game must satisfy $\pi_{11} < \frac{1}{2}$ and $\pi_{21} < \frac{1}{2}$. Therefore, the principal branch cannot be converging to the mixed-strategy equilibrium, and so must converge to the risk-dominant Nash equilibrium. \square

The conditions of the preceding theorem are satisfied by the logistic specification. Any function f_j derived from a probability distribution that is admissible in the sense of McKelvey and Palfrey will satisfy condition (2) of the theorem. Condition (3) does not follow directly even if f_j is derived from an admissible distribution, though it is a reasonable condition if λ is interpreted as a precision parameter.

This result is related to the results in ANDERSON, GOEREE, AND HOLT [2], who show that logit equilibria maximize a stochastic potential function, and apply the result to minimum-effort coordination games. For the case of 2×2 games, however, the logit specification is not needed to select the risk-dominant equilibrium. Rather, the selection of the risk-dominant equilibrium is a natural consequence of quantal response ideas. This implies that the usefulness of the logit solution concept as an equilibrium selection method in laboratory games may be limited, as risk-dominance does not appear to be the salient feature of games of this sort. (See STRAUB [18] and VAN HUYCK ET AL [20] for laboratory studies of coordination games comparing risk-dominance and payoff-dominance.)

7 Conclusions

This paper has presented a technique for efficiently tracing a branch of the logit equilibrium correspondence, with application to the problem of computing a single Nash equilib-

rium. The path followed by a branch of the correspondence is interpreted in terms of the replicator dynamics and best response dynamics, analogous to the interpretation of the Tracing Procedure as an adjustment between prior and final beliefs in an introspective process.

The presentation in this paper focused primarily on the principal branch. The path-following aspects of the analysis hold for all branches of the correspondence; the method can be applied to any branch, once a point on the branch has been identified (for example, by a grid search). Since generically the branches of the logit equilibrium correspondence other than the principal branch connect pairs of Nash equilibria, a modification of the procedure outlined can be used to compute another Nash equilibrium, given one is known. The technical problem that such a homotopy needs to start at parameter $\lambda = \infty$ can be finessed by restating the homotopy using parameter $\nu = \frac{\lambda}{1+\lambda}$.⁷ Thus, as λ ranges from 0 to ∞ , ν ranges from 0 to 1. At $\nu = 1$, the necessary starting point for such an attempt, the matrices involved in computing the homotopy are singular. However, Theorem 5 indicates that the branch is well-approximated by an initial starting direction involving only changes in ν . The author’s experience is this is feasible in many games, so long as care is given to the choice of initial step size so as to be sure that the matrices in (9) are suitably well-conditioned at the initial guess for a logit equilibrium for $\nu < 1$.

The characterization of the logit equilibrium correspondence branches, as well as the efficiency with which these branches can be traversed, suggests further investigation of logit equilibrium properties. Logit equilibrium ideas have been applied with some success in explaining deviations from Nash equilibrium predictions in laboratory games; see for example GOEREE AND HOLT [4] and references, and GOEREE, HOLT, AND PALFREY [5].

7. The parameterization of the quantal response correspondence by the “precision” parameter λ has been called the “West Coast” parameterization. The alternate specification of $\mu = \frac{1}{\lambda}$, where μ is a “noise” parameter, has been called the “East Coast” version. This parameterization by ν , then, is the “Texas” version; after all, Texas needs its own version of everything.

Standard techniques in numerical path following permit identification of points on a path where a given test function equals zero. These are generally used to detect branch or turning points in a path (in the case of this procedure, by using $|G_{-\lambda}|$ as the test function). For the purpose of maximum likelihood estimation, the techniques can be applied using the (directional) derivative of the likelihood function as the test function.

For many applications, joint estimates of λ and additional game parameters are desired. For example, Goeree, Holt, and Palfrey simultaneously estimate λ and a risk-aversion parameter. There is no immediate extension of the homotopy methods used here to trace out the surfaces of logit equilibria determined by multiple parameters. However, an adaptive search strategy in the dimensions of the additional parameters coupled by efficient identification of the maximum-likelihood λ for a given set of other parameters will yield improved results over a simple grid search.

The implementation of path-following techniques relies heavily on linear algebra. As such, it is possible to program the algorithm to exploit multicomputer environments, including Beowulf clusters. Although the complexity of truly large games will overwhelm even powerful systems, parallel implementations can still greatly extend the range of games of interest which can be analyzed.

Theorem 7 suggests an extension of risk-dominance beyond 2×2 games using quantal response ideas. The argument establishing Theorem 7 is based on monotonicity of the principal branch of the logit equilibrium correspondence. Using the dynamical interpretation of Section 4, at the centroid, each player prefers to play his strategy from the risk-dominant equilibrium; as λ increases, the second-order information about the other player's behavior reinforces this. An equilibrium at the end of a principal branch which is monotonic is then a natural candidate for being risk-dominant in an extended sense.

HOPKINS [10] points out that logit equilibria can be viewed as the stationary points of a stochastic best-reply dynamics. The eigenvalues of the matrix $G_{-\lambda}$, evaluated at a logit equilibrium, give information about the stability of the logit equilibrium under these

dynamics; additionally, bifurcations in the graph of the correspondence indicate changes in stability properties of the logit equilibria. The behavioral relevance of stability properties of logit equilibria, if any, remains uninvestigated.

Acknowledgments

A previous version of this paper was circulated and presented under the title “Computation of the Logistic Quantal Response Equilibrium Correspondence.” This work has benefited from discussions with John Dickhaut, Richard McKelvey, Tom Palfrey, Rajiv Sarin, John Van Huyck, and seminar participants at Caltech, Texas A&M, and at the Economic Science Association meetings in Tucson AZ, November 2001. The suggestions of two anonymous referees led to improved exposition of the results. David Kisel provided helpful assistance in preparing the final manuscript. All errors are the author’s.

Appendix A. Extensive games with perfect recall

MCKELVEY AND PALFREY [16] develop an agent quantal response equilibrium concept for extensive form games. Players choose their actions optimally at each information set, given what actions are chosen at all other information sets (including potentially subsequent information sets belonging to the player), and the random noise added to payoffs. The noise terms are taken to be independent across information sets of the same player. The path-following method described in the body applies directly to this setting as well; this appendix details the quantities that need to be calculated.

In what follows, β denotes a behavior strategy profile, i and l index players, j and m index information sets, and a and b denote actions. At each information set h_{ij} of player i , he has available a set of actions denoted by $A(h_{ij})$. Payoffs are now calculated conditional upon reaching an information set; $\bar{u}_{ija}(\beta)$ is defined as the expected payoff, conditional on reaching the information set h_{ij} , to player i of playing action $a \in A(h_{ij})$, when β is played

elsewhere in the tree. Since for finite λ the logit quantal response equilibria are on the interior of the set of behavior strategy profiles, the probability of reaching the information set is positive, and the conditional expectation is well-defined. McKelvey and Palfrey define a behavior strategy β to be a logistic agent quantal response equilibrium if

$$\beta_{ija} = \frac{e^{\lambda \bar{u}_{ija}(\beta)}}{\sum_{a' \in A(h_{ij})} e^{\lambda \bar{u}_{ija'}(\beta)}}.$$

The elements of the Jacobian for the homotopy are

$$\frac{\partial H_{ija}(\beta)}{\partial \beta_{lmb}} = \sum_{a' \in A(h_{ij})} \lambda \frac{\partial \bar{u}_{ija'}(\beta)}{\partial \beta_{lmb}} e^{\lambda \bar{u}_{ija'}(\beta)} - \lambda \frac{\partial \bar{u}_{ija}(\beta)}{\partial \beta_{lmb}} e^{\lambda \bar{u}_{ija}(\beta)},$$

for all actions b such that $b \in h_{lm}$ for all pairs $(l, m) \neq (i, j)$, and

$$\frac{\partial H_{ija}(\beta)}{\partial \lambda} = \beta_{ija} \sum_{a' \in A(h_{ij})} \bar{u}_{ija'}(\beta) e^{\lambda \bar{u}_{ija'}(\beta)} - \bar{u}_{ija}(\beta) e^{\lambda \bar{u}_{ija}(\beta)}.$$

These are, up to a change in notation, equivalent to the expressions obtained for the normal form case. The substitutions used in obtaining (9) can then be repeated analogously in the extensive form case.

All that remains is to consider the computation of the quantity $\frac{\partial \bar{u}_{ija}(\beta)}{\partial \beta_{lmb}}$. If the extensive form game has perfect recall, computation of this derivative can be done using computer code likely already written for the game. Write $N(h)$ to be the set of nodes belonging to information set h . Let $\bar{u}_{ija|n}(\beta)$ denote the payoff to playing action a at a node $n \in h_{ij}$. Denote by $P_n(\beta)$ the probability of reaching a node n given profile β , and $P_h(\beta)$ the probability of reaching an information set h . Then

$$\begin{aligned} \bar{u}_{ija}(\beta) &= \sum_{n \in N(h)} P_{n|h}(\beta) \bar{u}_{ija|n}(\beta) \\ &= \sum_{n \in N(h)} \frac{P_n(\beta)}{P_h(\beta)} \bar{u}_{ija|n}(\beta) \\ \bar{u}_{ija}(\beta) P_h(\beta) &= \sum_{n \in N(h)} P_n(\beta) \bar{u}_{ija|n}(\beta) \end{aligned}$$

$$\begin{aligned}
\sum_{n \in N(h)} P_n(\beta) \bar{u}_{ija|h}(\beta) &= \sum_{n \in N(h)} P_n(\beta) \bar{u}_{ija|n}(\beta) \\
\sum_{n \in N(h)} \frac{\partial P_n(\beta)}{\partial \beta_{lmb}} \bar{u}_{ija|h}(\beta) + P_n(\beta) \frac{\partial \bar{u}_{ija|h}(\beta)}{\partial \beta_{lmb}} &= \sum_{n \in N(h)} \frac{\partial P_n(\beta)}{\partial \beta_{lmb}} \bar{u}_{ija|n}(\beta) + \\
&\quad \sum_{n \in N(h)} P_n(\beta) \frac{\partial \bar{u}_{ija|n}(\beta)}{\partial \beta_{lmb}} \\
P_h(\beta) \frac{\partial \bar{u}_{ija}(\beta)}{\partial \beta_{lmb}} &= \sum_{n \in N(h)} \frac{\partial P_n(\beta)}{\partial \beta_{lmb}} [\bar{u}_{ija|n}(\beta) - \bar{u}_{ija}(\beta)] + \\
&\quad \sum_{n \in N(h)} P_n(\beta) \frac{\partial \bar{u}_{ija|n}(\beta)}{\partial \beta_{lmb}}.
\end{aligned}$$

When the game has perfect recall, probability of reaching a node n is the product of the action probabilities along the path reaching the node:

$$P_n(\beta) = \prod_{a \prec n} \beta_a.$$

Therefore,

$$\frac{\partial P_n(\beta)}{\partial \beta_{lmb}} = \prod_{a \prec n, a \neq b} \beta_a,$$

which is the same as the probability of reaching node n when b is played with probability one, and all actions at other information sets are played as specified in β . A similar observation applies to the quantity $\frac{\partial \bar{u}_{ija|n}(\beta)}{\partial \beta_{lmb}}$. Therefore, the payoff derivatives needed to compute the homotopy path can be computed without additional computer code.

Bibliography

- [1] E. L. Allgower and K. Georg. *Numerical Continuation Methods: An Introduction*. Springer-Verlag, Berlin, 1990.

- [2] S. P. Anderson, J. K. Goeree, and C. A. Holt. Minimum-effort coordination games: Stochastic potential and logit equilibrium. *Games and Economic Behavior*, 34:177--199, 2001.
- [3] C. D. Garcia and W. I. Zangwill. *Pathways to Solutions, Fixed Points, and Equilibria*. Prentice-Hall, Englewood Cliffs NJ, 1981.
- [4] J. K. Goeree and C. A. Holt. Stochastic game theory: For playing games, not just for doing theory. *Proceedings of the National Academy of Sciences*, 96:10564--10567, 1999.
- [5] J. K. Goeree, C. A. Holt, and T. Palfrey. Quantal response equilibrium and overbidding in private value auctions. *Journal of Economic Theory*, 104:247--272, 2002.
- [6] S. Govindan and R. Wilson. A global Newton method to compute Nash equilibria. *Journal of Economic Theory*, 110:65--86, 2003.
- [7] S. Govindan and R. Wilson. Computing Nash equilibria by iterated polymatrix approximation. *Journal of Economic Dynamics and Control*, in press.
- [8] J. Harsanyi and R. Selten. *A General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge MA, 1988.
- [9] P. J.-J. Herings and A. van den Elzen. Computation of the Nash equilibrium selected by the tracing procedure in n -person games. *Games and Economic Behavior*, 38:89--117, 2002.
- [10] E. Hopkins. Two competing models of how people learn in games. *Econometrica*, 70:2141--2166, 2002.
- [11] K. L. Judd. *Numerical Methods for Economics*. MIT Press, Cambridge MA, 1998.
- [12] E. Kohlberg and J.-F. Mertens. On the strategic stability of equilibria. *Econometrica*, 54:1003--1038, 1986.
- [13] C. E. Lemke and J. T. Howson. Equilibrium points of bimatrix games. *Journal of the Society of Industrial and Applied Mathematics*, 12:413--423, 1964.
- [14] R. D. Luce. *Individual Choice Behavior: A Theoretical Analysis*. John Wiley & Sons, Inc., New York, 1959.
- [15] R. D. McKelvey and T. R. Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10:6--38, 1995.

- [16] R. D. McKelvey and T. R. Palfrey. Quantal response equilibria for extensive form games. *Experimental Economics*, 1:9--41, 1998.
- [17] R. B. Myerson. Refinements of the Nash equilibrium concept. *International Journal of Game Theory*, 7:73--80, 1978.
- [18] Paul G. Straub. Risk dominance and coordination failures in static games. *Quarterly Review of Economics and Finance*, 34(4):339--363, 1995.
- [19] P. Taylor and L. Jonker. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 40:145--156, 1978.
- [20] John B. Van Huyck, Joseph B. Cook, and Raymond C. Battalio. Adaptive behavior and coordination failure. *Journal of Economic Behavior and Organization*, 32:483--503, 1997.
- [21] R. Wilson. Computing simply stable equilibria. *Econometrica*, 60:1039--1070, 1992.
- [22] Y. Yamamoto. A path-following procedure to find a proper equilibrium of finite games. *International Journal of Game Theory*, 22:249--259, 1993.