

Chapter 3. QRE Explanations of Intuitive Behavioral Anomalies

Probabilistic choice models, for both individual discrete choices and for games between two or more players, are motivated by empirical studies in which observed decisions exhibit some noise. The empirical applications to be considered in this chapter are based on laboratory experiments that implement games with players who are motivated by cash payoffs. Anyone who has looked at decisions made by human subjects in such situations will know that, despite the strong empirical regularities, there is often a fair amount of unexplained variation across individuals and over time for the same individual. Some of the noise is due to recording errors. For example, we have observed subjects who are shown a signal (e.g. draw of a colored marble from a cup) to occasionally write down the wrong signal. Therefore, it is easy to imagine that decisions are sometimes recorded incorrectly, especially in low payoff situations. Sometimes the variations seem to be due to emotional reactions, e.g. to others' prior decisions, or to individual differences in attitudes about relative payoffs or risk. The latter factors can, of course be incorporated into the models, as will be seen in chapter 4, but there will surely be other factors that are unobserved or unmodeled, and hence, that constitute noise effects. Regardless of the source, players in games will respond systematically to the anticipated behavior of other players, which should include noise effects to the extent that they are present. The first section provides some notation and background for the logit equilibrium analysis of symmetric games to be covered in this chapter.

A well-designed experiment will often involve a comparison of behavior in different games obtained by altering the number of players in the game or some parameter of the function that determines money payoffs. For example, it is frequently possible to change a parameter that has no effect on the Nash equilibrium, but that produces an imbalance between the cost of deviations in one direction relative to the cost of deviations in another direction. Recall that changes in relative payoffs alter the choice probabilities obtained from the response functions considered in Chapter 2, and as a result, one would expect to observe more deviations in the direction with lower deviation costs.

We will begin by revisiting several variations of the asymmetric matching-pennies game that was used extensively in the previous chapter. In a mixed-strategy Nash equilibrium for this game, each player is indifferent between the two decisions, since otherwise perfect rationality would require playing the better strategy with probability 1. It was shown that if a symmetric version of this game is altered by making one player's payoff much higher or lower in a particular cell, the player's Nash equilibrium choice probability for this decision should be unchanged, as required to keep the *other* player indifferent between decisions over which randomization occurs. The second section of this chapter reports data from laboratory experiments that were structured to test for the absence of an "own-payoff effect." The logit quantal response equilibrium is shown to generate such own-payoff effects by "smoothing" off the sharp corners of the Nash best-response functions and making them "flatter" in the sense

of being less responsive to payoff differences.

The third section pertains to a class of n -person games with two decisions, e.g. whether or not to vote, volunteer, or contribute to a public good. For example, Kahneman (1988) reported an early experiment in which subjects had to choose whether or not to enter a market, and the observed number of entrants approximately equal to a market capacity parameter. He remarked “To a psychologist, it looks like magic.” This tendency for aggregate behavior to equate expected payoff levels for entry and exit (as predicted in a Nash equilibrium) has been documented by others (e.g. Ochs, 1990; Meyer *et al.*, 1992). Some subsequent experimental studies, however, report over-entry and others report under-entry, and some even find over-entry relative to Nash predictions in some treatments and under-entry in others. We show that the addition of “noise” in these models tends to flatten the best response functions in the relevant range, which generates explanations of behavioral anomalies such as the divergence of results on whether there is excess entry or not. Specific applications involving binary choice games are considered next: the “volunteer’s dilemma” in section 3.4, voting participation games in section 3.5, and jury voting in 3.6. In these sections, the focus is on changes in voting rules and/or group sizes on Nash and QRE predictions, which will be compared with data from laboratory experiments.

In games with decisions that are constrained to be in some interval of numerical values, the distribution of decisions in a quantal response equilibrium is not merely a (possibly asymmetric) spread to each side of a Nash equilibrium, since “feedback effects” from deviations by one player alter others’ expected payoff profiles, which would induce further changes. The Traveler’s Dilemma experiment, considered section 3.7, provides an example in which these feedback effects can cause the quantal response equilibrium distribution of decisions to differ dramatically from the unique Nash equilibrium.¹

The Traveler’s Dilemma has the interesting property that payoffs are determined by the *minimum* of players’ decisions, a property that arises naturally in many economic interactions. For example, in a market for a relatively homogenous product, the firm with the lowest price would be expected to obtain a larger market share. Section 3.8 pertains to a simple price-competition game, which is again motivated by the possibility of changing a payoff parameter that has no effect on the unique Nash equilibrium, but which may be expected to affect the quantal response equilibrium. In particular, it is intuitive to expect that a reduction in the market share advantage enjoyed by the low-price firm would raise the observed distribution of prices.

In many production settings, each player’s contribution to a joint effort is critical in the sense that a reduction in one player’s effort causes a bottleneck that reduces the payoffs to all. In this case, it is the minimum of all players’ efforts that determines payoffs, along with the cost of a player’s own effort. This is a “coordination game” when the payoff function is such that nobody would wish to reduce the minimum effort that determines the joint production), but a unilateral increase from a common effort level is individually costly and

¹The reader may wish to play the Traveler’s Dilemma before reading further. This can be done online by going to the site: <http://veconlab.econ.virginia.edu/tddemo.htm>. This demo takes about ten minutes; it consists of five rounds of play in which you will be matched with stored data taken from a University of Virginia Law School class.

has no effect on the minimum. In the Minimum-Effort Coordination Game to be analyzed in the section 3.9, any common effort in the range of feasible effort levels is a Nash equilibrium, but one would expect that an increase in the cost of individual effort or an increase in the number of players who are trying to coordinate would reduce the effort levels observed in an experiment.

Many non-price allocations are made in response to lobbying and other “rent-seeking” activities, and there is a fair amount of evidence that the social costs of such activities may consume a large fraction of the value of the prize being allocated. The full dissipation of rent would result in negative expected payoffs, which can never occur in a Nash equilibrium with a zero-profit no-loss opportunity. However, full dissipation (and even over dissipation) has been observed in laboratory experiments, especially with large numbers of contestants. The eighth section presents an analysis of the logit equilibrium and rent dissipation for a rent-seeking contest that is modeled as an “all-pay auction.”

This chapter will consider the Nash and quantal response equilibrium predictions for these and other applications. In each case, we discuss results of laboratory experiments with two or more treatments that were designed to provide a sharp focus on differences between Nash and QRE predictions.

1. The Logit Equilibrium

Most of the analysis in this chapter will involve a comparison of data from laboratory experiments with the predictions of the logit quantal response equilibrium that was introduced in Chapter 2. Before considering specific games, it is useful to begin with a little review and background, presented in the context of the special case of symmetric two-person games, which simplifies the notation and clarifies the intuition. Recall that there are three steps to be taken to characterize a logit equilibrium: 1) calculating the expected payoffs for each decision, given belief probabilities, 2) using these expected payoffs with the response functions to generate the choice probabilities, and 3) imposing the equilibrium condition that the belief probabilities match the corresponding choice probabilities.

In the symmetric games to be considered, each player has the same number of strategies, J , which are denoted: x_1, x_2, \dots, x_J , where the player subscripts have been suppressed. First, consider the payoff function for a player who selects decision x_j . In a two-person game, the actual payoff depends on the other player’s decision, k , and will be denoted $v(x_j, x_k)$. If this person is a Row player in a two-person matrix game, for example, this would be the entry in the j th row and the k th column. Let p_k denote the probability representing a player’s beliefs that the other will choose decision k . Then the expected payoff for selecting decision j , denoted π_j , is calculated in a straightforward manner:

$$\text{Payoff Functions: } \pi_j = \sum_k v(x_j, x_k) p_k, \quad j = 1, \dots, J. \quad (1.1)$$

A regular quantal response function would determine the choice probabilities for each decision

as increasing functions of the expected payoff for that decision. In a logit model, the choice probabilities are specified to be ratios of exponential functions:

$$\text{Response Functions: } \sigma_j = \frac{e^{\lambda\pi_j}}{\sum_k e^{\lambda\pi_k}}, \quad j = 1, \dots, J, \quad (1.2)$$

where λ is a precision parameter that determines the degree of rationality in the model, ranging from complete randomness ($\lambda = 0$) to full rationality ($\lambda = \infty$). In equilibrium, the belief probabilities, p_j , used to calculate the expected payoffs in (3.1.1) must match the choice probabilities, σ_j , determined by the response function in (3.1.2):

$$\text{Equilibrium Conditions: } p_j = \sigma_j, \quad j = 1, \dots, J. \quad (1.3)$$

In this probabilistic sense, the logit quantal response equilibrium requires a consistency of actions and beliefs. For a specific value of the precision parameter, λ , equations (3.1.1) and (3.1.3) can be substituted into (3.1.2) to obtain J equations in the equilibrium mixed strategy probabilities, σ_j . These equations are then solved simultaneously to find the logit QRE.

Even though the quantal response equilibrium involves mixed strategies, due to the probabilistic nature of the response functions, it is typically quite different from a mixed-strategy *Nash* equilibrium, which involves randomization by perfectly rational players. In particular, a perfectly rational player would not be willing to randomize over a set of decisions unless the expected payoffs for all decisions were exactly equal. In contrast, the equilibrium expected payoffs for different decisions are typically not equal in a quantal response equilibrium, which puts more probability mass on the payoffs with higher expected payoffs.

With only two decisions, the equilibrium condition for each person is:

$$\sigma_1 = \frac{e^{\lambda\pi_1}}{e^{\lambda\pi_1} + e^{\lambda\pi_2}} = \frac{1}{1 + e^{\lambda(\pi_2 - \pi_1)}}. \quad (1.4)$$

The intuition behind the effects of changes in the precision parameter can be seen from (3.1.4). As the amount of noise is increased, so that $\lambda \rightarrow 0$, the equilibrium probabilities converge to $1/2$, regardless of how the expected large payoff differences happen to be. In the other extreme, as $\lambda \rightarrow \infty$, the term $\lambda(\pi_2 - \pi_1)$ goes to infinity if decision two has a higher expected payoff throughout this range, so σ_1 goes to 0. Conversely, σ_1 goes to 1 if decision 1 has the higher expected payoff, as required under perfect rationality. Thus the two extreme limits of the λ parameter generate the two extremes of perfectly noisy and perfectly rational behavior. For intermediate cases, it is clear from the final term in (3.1.4) that an increase in the precision λ will make the logit function more responsive, i.e. with sharper "corners."

As noted in Chapter 2, there are two alternative motivations for the specification of the quantal response function: *structural* and *reduced-form*. The structural interpretation of the logit model is based on an assumption that each of the decisions has a payoff that is perturbed by a random shock, so that decision 1 is chosen if $\pi_1 + \varepsilon_1 > \pi_2 + \varepsilon_2$, where the shocks have an extreme value distribution on the real line with a distribution function $G(\varepsilon_i) = \exp(-e^{-\lambda\varepsilon_i})$ and density $g(\varepsilon_i) = \exp(-e^{-\lambda\varepsilon_i})\lambda e^{-\lambda\varepsilon_i}$. For a given value of ε_1 , decision 1 is chosen with

probability $G(v_1 - v_2 + \varepsilon_1)$, which can be integrated using the density of ε_1 : $\Pr\{\text{decision 1}\} = \int_{-\infty}^{\infty} G(v_1 - v_2 + \varepsilon_1) f(\varepsilon_1) d\varepsilon_1$. Then a transformation of variables can be used to express this integral as the logit response in equation (3.1.4).²

Another perspective on the logit model is that it provides a reduced-form response function with interesting properties. One such property is that the ratio of choice probabilities depends on the expected payoffs for those decisions but is independent of the payoffs for the other decisions that are “irrelevant” to the comparison:

$$\text{Independence of Irrelevant Alternatives: } \frac{\sigma_i}{\sigma_j} = \frac{e^{\lambda\pi_i}}{e^{\lambda\pi_j}}, \quad \text{for all } i, j. \quad (1.5)$$

A second property of the logit response function is that the addition of a constant, w to each of the J expected payoffs in the numerator and denominator of (3.1.2) will produce $e^{\lambda w}$ terms that will all cancel, so the addition of a constant to all payoffs will not affect the choice probabilities. As Luce (1959) noted, the reverse logic also applies, i.e. any reduced-form quantal response function with these properties (independence of irrelevant alternatives and invariance with respect to additive payoff increments) must have the logit form.³ These properties have a certain appeal, but they can be controversial in some situations.⁴

2. Asymmetric Matching Pennies

Table 3.1 Shows a matching pennies game with the same general structure as that considered in the previous chapter. If the Row player’s payoff parameter, X , is set equal to 72,

²Note that: $\Pr\{\text{decision 1}\} = \int_{-\infty}^{\infty} \exp(-e^{-\lambda[v_1 - v_2 + \varepsilon_1]}) \exp(-e^{-\lambda\varepsilon_1}) \lambda e^{-\lambda\varepsilon_1} d\varepsilon_1$. To integrate this, use a change of variables with $\tau = e^{-\lambda\varepsilon_1}$ and $d\tau = -\lambda e^{-\lambda\varepsilon_1} d\varepsilon_1$, where $\tau \in (-\infty, 0)$. The resulting integral is: $\int_{-\infty}^0 \exp(-e^{-\lambda v_1} e^{\lambda v_2} \tau) \exp(-\tau) (-d\tau)$. Using a sign change to reverse the limits of integration, this integral can be expressed: $\int_0^{\infty} \exp(-\tau[1 + e^{\lambda v_2}/e^{\lambda v_1}]) d\tau$, which integrates to yield (3.1.4).

³To see this, suppose that the choice probability for each decision is a function of its expected payoff and of the expected payoffs of all other decisions. Without loss of generality, these probabilities can be written as: $\sigma_j = f(\pi_j)g_j(\pi_1, \dots, \pi_J)$, where the $f(\cdot)$ functions are the same for all decisions (no j subscript). With this notation, the ratios of choice probabilities can be written: $\sigma_i/\sigma_j = (f(\pi_i)/f(\pi_j)) \cdot (g_i/g_j)$ for all i, j . The independence assumption implies that the g_i/g_j ratios are equal to constants, k_{ij} , which are independent of the expected payoffs, even though each of the $g_i(\cdot)$ functions will generally depend on all expected payoffs. Thus $\sigma_i/\sigma_j = (f(\pi_i)/f(\pi_j))k_{ij}$. Next, the independence with respect to additive payoff shifts implies that: $\sigma_i/\sigma_j = (f(\pi_i)/f(\pi_j))k_{ij} = (f(\pi_i + w)/f(\pi_j + w))k_{ij}$ for all i, j and all w . The final equation can be expressed as $f(\pi_i)f(\pi_j + w) = f(\pi_j)f(\pi_i + w)$, which when differentiated with respect to w and evaluated at $w = 0$ yields: $f'(\pi_j)/f(\pi_j) = f'(\pi_i)/f(\pi_i)$ for all i, j . Thus, the “growth rate” of the $f(\cdot)$ function is a constant, λ , and hence the function is an exponential of the form $e^{\lambda\pi_i}$.

⁴For example, a comparison between alternatives A and B may be affected by whether there is a third (“irrelevant”) alternative, C , that is dominated in all dimensions by A , which may increase the chances that A will be selected. Moreover, an expected-payoff difference of a hundred dollars might have a large impact on decisions for small consumer durable purchases, but it may not affect home purchases, where the un-modeled random effects might dwarf minor cash differences. In situations with large changes in the payoff scale or when payoffs are quite variable, it may be advisable to consider alternatives to the logit specification, e.g. the power function response function introduced in chapter 2.

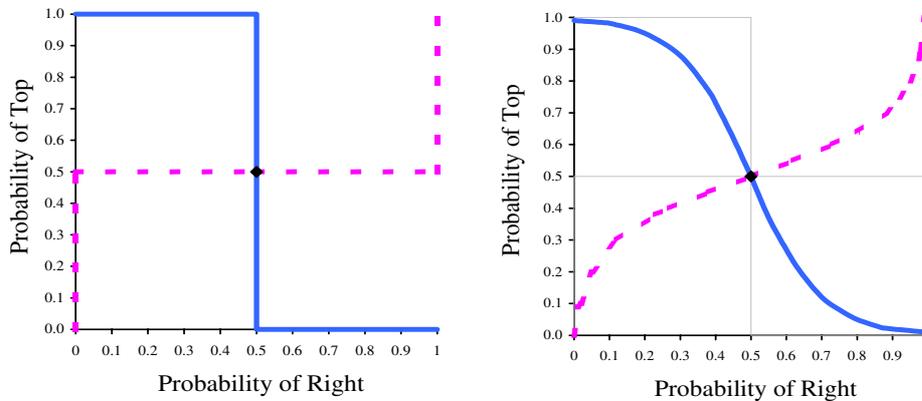


Figure 1: *Response functions for a symmetric matching-pennies game with $\lambda = \infty$ on the left and $\lambda = 0.1$ on the right.*

then the payoffs are symmetric in the sense that both players would be indifferent between their two decisions if the other player randomizes with probability $1/2$.

Table 3.1: *Generalized matching pennies games ($X = 40, 72, \text{ or } 360$).*

	Left	Right
Top	$X, 36$	$36, 72$
Bottom	$36, 72$	$72, 36$

The panel on the right side of Figure 3.1 shows logit response functions for this game for $\lambda = 0.1$, where the probability of Row's choice of Top is measured vertically and the probability of Column's choice of Right is measured horizontally. (Note that an increase in the probability of right corresponds to a movement to the right, which is different from the convention used in the previous chapter.) These response functions are determined by (3.1.2). In particular, let the choice probabilities for Row's choice of Top and Column's choice of Left be denoted by σ_T and σ_L respectively, where the player subscripts have been suppressed. Similarly, let the expected payoffs be denoted by π_T and π_B for Row and by π_L and π_R for Column. In equilibrium, the choice probabilities for one person will match the probabilities that represent the other's beliefs, so we can calculate Row's expected payoffs as:

$$\pi_T = X(1 - \sigma_R) + 36\sigma_R \quad \text{and} \quad \pi_B = 36(1 - \sigma_R) + 72\sigma_R, \quad (2.1)$$

which are substituted into (3.1.4) to determine Row's choice probability, σ_T as a function of

the probability σ_R that Column chooses Right. For $\lambda = 0.1$, this response function is shown as a solid line in the right panel of Figure 3.1. Similarly, Column's expected payoffs are:

$$\pi_L = 36\sigma_T + 72(1 - \sigma_T) \quad \text{and} \quad \pi_R = 72\sigma_T + 36(1 - \sigma_T), \quad (2.2)$$

which can be used to determine Column's response as a function of σ_T , which is shown as a dashed line on the right side of Figure 3.1.

In the limit at $\lambda \rightarrow \infty$, the dashed best response line in the left panel starts on the left side when Row's probability of playing Top is less than $1/2$, and it crosses over at $1/2$, since Column's best response is to play Right when Row is more likely to play Top. Row's best response line, shown as a solid line, is constructed in a similar manner. The intersection of these lines, at probabilities of $1/2$, corresponds to the mixed-strategy Nash equilibrium. The introduction of some noise tends to smooth these functions, as shown in the right panel of Figure 3.1, which was drawn for $\lambda = 0.1$. Note that, in this symmetric-payoff setup, there is no equilibrium effect of noise on the intersection of the response function. Hence, it is not surprising that experiments done with symmetric structures tend to yield choice proportions that are close to the Nash (and QRE) prediction of $1/2$ (e.g., Ochs, 1995; Goeree and Holt, 2001).

Next suppose that the value of Row's payoff parameter X is reduced from 72 to 40 in Table 3.1. This reduction makes the Top strategy a lot less attractive, so Row's best response line (shown as the "sharp" thin line in the left panel of Figure 3.2) will drop to the bottom of the figure as soon as Column's probability of Right is above 0.10. The intersection of the best response functions in the left panel determines the Nash mixed-strategy equilibrium at $(\sigma_R = 0.10, \sigma_T = 0.50)$, as indicated by the intersection of the thin lines. The logit equilibrium with $\lambda = 0.1$ is shown as the intersection of the thick curved (dashed and solid) lines in the left panel of the figure. In a Nash equilibrium, the Column player responds to the reduction in Row's payoff for Top by playing Right less often, and the equilibrium effect is for Row to continue playing Top with probability $1/2$, despite the large reduction in the (Top, Left) payoff. In contrast, the logit equilibrium predicts an "own-payoff effect," i.e. that Row will choose Top less often. The "40 data" point in the figure shows the choice proportions for an experiment done with this low value of X , with three groups of 10-12 subjects, who were randomly matched for 25 rounds.

The same subjects also played 25 rounds for a game with the asymmetry in the opposite direction, i.e. with an increase in Row's (Top, Left) payoff to 360, which raises Row's best response line to the top of the right panel of Figure 3.2 unless Column plays Right with a probability in excess of 0.9. This change moves the Nash equilibrium to the right horizontally, with Column's equilibrium adjustment to Row's payoff change neutralizing its effects, i.e. with no "own-payoff effect." In contrast, the data proportions do exhibit an own-payoff effect, as shown by the "360 data" dot on the right side of the figure. This effect is somewhat over-predicted by the intersection of the logit response lines for $\lambda = 0.1$, even though this parameter provides an excellent fit for the $X = 40$ treatment. One possible explanation for this prediction error is that the very high payoff of 360 is undervalued by subjects who are risk averse (Goeree, Holt, and Palfrey, 2003). Similar own-payoff effects

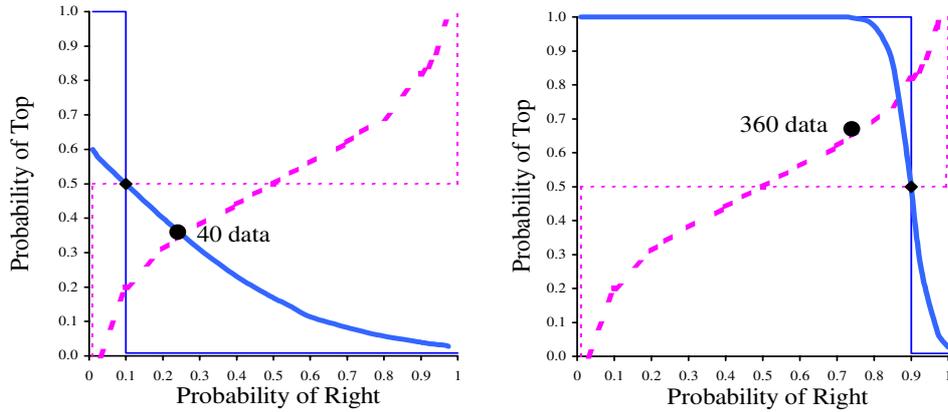


Figure 2: *Response functions for an asymmetric matching-pennies game with $X = 40$ on the left and $X = 360$ on the right.*

that contradict Nash predictions are reported by Ochs (1995) and McKelvey, Palfrey, and Weber (2000). In all cases, the QRE predictions track the qualitative features of the data, although the effects of large increases in the (Top, Left) payoff are under-predicted by the logit QRE model when risk neutrality is assumed.

The logit calculations thusfar have used expected money payoffs, but in principle, there is no reason not to use expected utilities. An empirical implementation would require the joint estimation of the logit precision parameter and one or more risk aversion parameters, which will be discussed in the next chapter.

3. Binary-Choice Participation Games

In many n -person games, players must choose between two decisions, e.g. whether or not to enter, vote, volunteer, or contribute to a public good. In this section, we consider a set of closely related binary-choice games, where the payoff for one of the decisions depends on the number, m , who make that decision ($m \leq n$). In a market-entry game, for example, the payoff for each entrant will generally be a decreasing function of the number of entrants, and the payoff for exit will be a non-decreasing function of m that represents the opportunity cost of entry.

A mixed strategy in this game will be characterized by a probability, p , of the active participation strategy, e.g., entry, voting, volunteering, contributing, etc., and $1 - p$ will denote the probability of the “exit-out” strategy. In a symmetric equilibrium in which all others participate with the same probability p , a player’s expected payoff from participation will depend on both p and the number of others ($n - 1$) and is assumed to be finite. When

participation choices are independent, the expected payoff for participating can be calculated from a binomial distribution with parameters $n - 1$ and p , and will be denoted $\pi_1(n - 1, p)$. Note that this is the expected payoff for a player who participates (*with probability 1*) when all $n - 1$ others participate with probability p . Similarly, the expected payoff from exit, denoted $\pi_2(n - 1, p)$, represents the payoff from entry with probability 0 when the others enter with probability p .

Let the random payoff shocks to the π_1 and π_2 payoffs be denoted by ε_1 and ε_2 respectively, which are assumed to be independently and identically distributed with mean zero and positive density on $(-\infty, \infty)$. The shocks are scaled by an error parameter, μ , and therefore, participation occurs if $\pi_1 + \mu\varepsilon_1 > \pi_2 + \mu\varepsilon_2$, or equivalently, if $\lambda(\pi_1 - \pi_2) > \varepsilon_2 - \varepsilon_1$, where the precision λ is the reciprocal of the μ . Let $F(\cdot)$ denote the distribution function of the difference between the shocks, which is increasing on $(-\infty, \infty)$. It follows that the equilibrium probability of participation must satisfy:

$$p = F\left(\lambda[\pi_1(n - 1, p) - \pi_2(n - 1, p)]\right). \quad (3.1)$$

When μ is large (λ is small), the effects of the shocks are magnified, and expected payoff differences matter less. Expected payoff maximization occurs in the other extreme for low values of the error parameter (high precision). Equation (3.3.1) characterizes a quantal response equilibrium, since the probability used to calculate expected payoffs on the right matches the response probability on the left.

By applying the inverse of the distribution function to both sides of equation (3.3.1), one obtains an expression that separates the factors affecting the noise terms, which are on the left, from the expected payoff difference on the right:

$$F^{-1}(p)/\lambda = \pi_1(n - 1, p) - \pi_2(n - 1, p). \quad (3.2)$$

Before proceeding, it will be useful to derive some general properties of the inverse distribution function that do not depend on any particular parametric specification. It follows from the i.i.d. assumption that $\Pr\{\varepsilon_2 < \varepsilon_1\} = 1/2$, so $\Pr\{\varepsilon_2 - \varepsilon_1 < 0\} = 1/2$, and hence $F(0) = 1/2$. These properties of the distribution function imply that its inverse is increasing on $(0,1)$ with $F^{-1}(1/2) = 0$, as shown by the curved “inverse distribution” line in Figure 3.3.

Equation (3.3.2) implies that the quantal response equilibrium occurs at the point where the expected payoff difference line crosses $F^{-1}(p)/\lambda$, which will be called the *inverse distribution line*, even though it includes the effects of both λ and the inverse of $F(\cdot)$. Note that an increase in λ makes the inverse distribution line flatter in the center of Figure 3.3, and conversely, a reduction in precision (more noise) makes the line steeper in the center.

There are two expected payoff difference lines in Figure 3.3: a low payoff difference line on the left, and a higher payoff difference line, the relevant part of which is shown on the right. Both lines are continuously decreasing in p , as would be the case for games with congestion effects, i.e. in which participation by one person decreases the participation payoffs for others. For the game with the lower payoff difference line, the expected payoff for

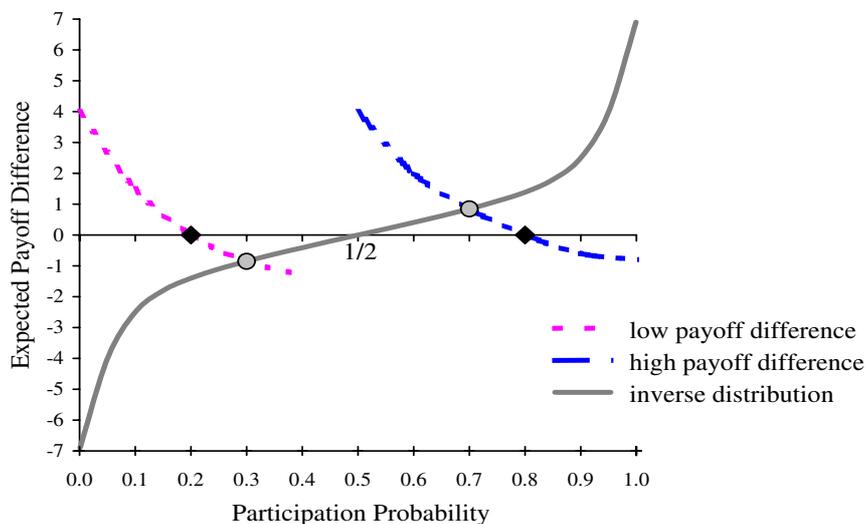


Figure 3: *Equilibrium points for an entry game with negative externalities*
 (Key: Nash equilibria are diamonds, QRE are circles.)

participation is higher than for exit if the $N - 1$ others participate with a probability that is less than 0.2, and the expected payoff for exit is higher to the right of 0.2 in the figure. When $p = 0.2$, the expected payoffs are equal, and this point represents a Nash equilibrium in mixed strategies. The Nash equilibrium for the game with the higher payoff difference line is located at 0.8, as shown by the diamond mark on the right side of the figure. The quantal response equilibrium for each case is marked with a large dot.

Since the inverse distribution line is continuously increasing with a range from $-\infty$ to ∞ , and the expected payoff difference line is decreasing for games with negative externalities, a symmetric quantal response equilibrium will exist and it will be unique. Moreover, any factor that shifts the expected payoff difference line downward will lower the QRE participation probability. For example, in a market entry game, an increase in the number of potential entrants will raise the expected number of entrants for each given level of p , which will reduce the expected payoff from entry as a function of p . If the cost of entry is not affected, then it follows that an increase in the number of potential entrants will lower the quantal response entry probability. The same logic can also be used to derive a similar prediction for the Nash probability of entry. One way in which the Nash and quantal response predictions differ when the expected payoff difference is strictly decreasing in p is that the QRE prediction is always closer to $1/2$, which reflects the fact that noise pulls the probabilities closer to the center, and this difference will be accentuated when there is more noise (lower λ), which makes the inverse distribution line steeper around $1/2$.

The qualitative nature of the QRE prediction, that entry rates will be pulled toward $1/2$ relative to the Nash prediction, is generally borne out in market entry experiments. First, note that if the expected payoff line crosses the zero point at $p = 1/2$, then it will also intersect the inverse distribution line at that point, so the Nash and quantal response equilibria coincide. Meyer *et al.* (1992) ran an experiment with groups of size 6 and parameters in their baseline treatment which caused the equilibrium number of entrants to be 3, so the Nash equilibrium probability of entry was $1/2$. The average number of entrants was not statistically different from 3 for the baseline sessions, even when the game is repeated for 60 periods (see their Tables 3 and 5). In contrast, results that deviate from the Nash prediction are reported in Camerer and Lovo (1999). They ran an experiment in which subjects had to choose whether or not to enter a market with an integer-valued “capacity” of c . Entrants were randomly ranked and the top c entrants were given shares of a \$50 pie, based on their rank. The exit payoff was set to 0. The capacities and group sizes (14-16) were selected so that the Nash entry probability was greater than or equal to $1/2$ in all cases. Under-entry was observed in all of the eight baseline sessions, which resulted in positive expected payoffs from entry (see their Table 4).⁵ Some of the most striking evidence in support of the the QRE prediction is found in Sundali, Rapoport, and Seale. Their design used an entry game with 10 different capacity parameters, which resulted in Nash equilibrium entry rates at about 10 more or less equally spaced probabilities on the interval from 0 to 1. On average, there was over-entry in all 5 cases with Nash predictions below $1/2$, and there was under entry in 4 of the 5 cases with Nash predictions above $1/2$, with the only exception being a treatment in which the Nash prediction was close to $1/2$.

In many games of interest, participation by one player may tend to increase, not decrease, the participation payoff of others. For example, consider a production process in which individuals decide whether or not to help in a group production activity in which each person’s productivity is an increasing function of the number of others who participate. The game can be structured so that it does not pay to participate unless enough others do so, and therefore, the expected payoff difference is negative for low p , positive for high p , and increasing in p . Figure 3.4. illustrates this case of positive externalities. Since both the expected payoff difference and the inverse distribution lines are increasing, it is possible to have multiple crossings, i.e. multiple quantal response equilibria. In this case, any parameter change that causes the expected payoff difference line to shift upward will raise the equilibrium participation probability for some of the equilibria and lower it for others.

The indeterminacy of comparative statics effects for participation games with positive externalities can be dealt with using an analysis of dynamic stability. Suppose that the participation probabilities tend to increase when the current participation probability, p is less than the stochastic best response to p : $p < F(\lambda(\pi_1(n-1, p) - \pi_2(n-1, p)))$, or equivalently, when $F^{-1}(p)\lambda < \pi_1(n-1, p) - \pi_2(n-1, p)$. Thus the participation

⁵Camerer and Lovo (1999) also report a second treatment in which subjects were told that entrants would be ranked on the basis performance on a sports trivia quiz. They observed over-entry in this case, which they attributed to “overconfidence” and neglecting that fact that other participants might think in the same manner “reference group neglect.”

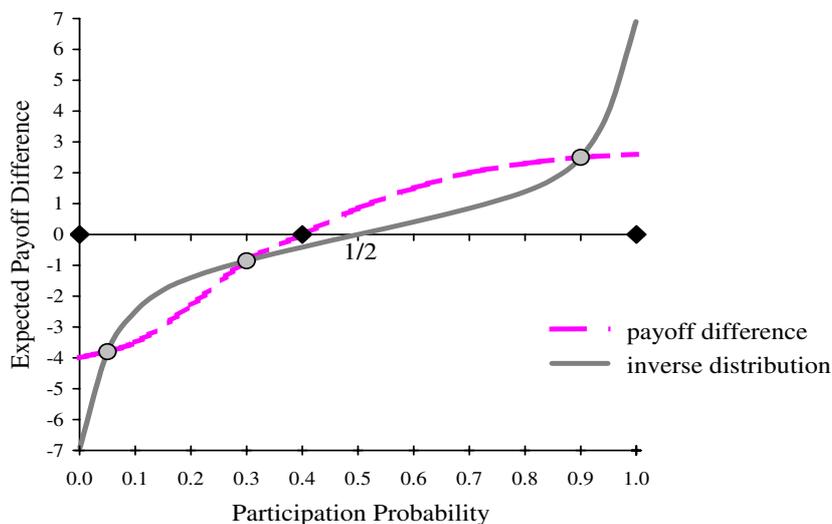


Figure 4: *Equilibrium points for an entry game with positive externalities*
 (Key: Nash equilibria are diamonds, QRE are circles.)

probability would tend to increase (decrease) when the inverse distribution line in the Figure 3.4 is below (above) the expected payoff difference line. For example, consider the quantal response equilibrium at $p = 0.3$ in the figure. At probabilities a little above 0.3, there is an upward pressure on the participation probability, and there is a downward pressure at probabilities a little below 0.3. Therefore this equilibrium is unstable. Conversely, the two more extreme quantal response equilibria, at probabilities of about 0.05 and 0.9, are stable. Note that the stable equilibria are those for which the inverse distribution line intersects the expected payoff difference line from below as one moves from left to right. For these stable equilibria, any parameter change that shifts the payoff difference line upward will raise the QRE probability of participation. One lesson to be drawn from this analysis is that in cases where comparative-statics effects are indeterminate or unintuitive, it may help to consider adjustment dynamics.

The results of this section can be summarized:

Proposition 3.1. There is at least one symmetric quantal response equilibrium in a symmetric binary-choice participation game. If the expected payoff difference line is a decreasing function of the probability of participation, then the QRE is unique and it lies between the Nash equilibrium probability and the probability of $1/2$. In this negative-externality case, and any exogenous factor that increases the participation payoff or lowers the exit payoff will raise the QRE participation probability. The same comparative statics result holds when there are multiple equilibria and attention is restricted to the ones that are stable.

4. The Volunteer's Dilemma

Many public goods games have the property that a person's decision to participate may benefit all other players. In some cases, a threshold level of participation is required for the group benefit to be obtained. This section pertains to a special case where the threshold is one, i.e. only a single volunteer is needed. For example, it only takes one person to cast a politically costly veto or to sponsor a bill authorizing a pay raise for all members of a legislative body. The dilemma here is that the act of volunteering is costly, and therefore, each person would prefer that someone else bear the cost.

The *volunteer's dilemma*, introduced by Diekmann (1986), is a game in which each of the n players in a group receives a benefit, B , if at least one of them incurs a cost, $C < B$. The expected payoff from volunteering is $B - C$, since a decision to volunteer ensures that the threshold is met. Next, consider the expected payoff from not volunteering in a symmetric equilibrium in which the $n - 1$ others volunteer with probability p . In this case, there is no cost, and there is a benefit B if at least one of the others volunteers, which is calculated as 1 minus the probability that none of them volunteer. Thus the expected payoff for not volunteering is $B(1 - (1 - p)^{n-1})$, and the expected payoff difference is $B(1 - p)^{n-1} - C$. This expected payoff difference function is decreasing in p , so this game is a special case of the entry game with negative externalities represented in Figure 3.4. Proposition 3.1 implies that the QRE probability is unique, decreasing in n and C , and increasing in B .

For comparison, consider the Nash equilibrium in mixed strategies, which is found by equating the expected payoff difference to 0 and solving for the Nash probability, p^* :

$$p^* = 1 - \left(\frac{C}{B}\right)^{\frac{1}{n-1}}. \quad (4.1)$$

This probability is increasing in B and decreasing in C and n . These results are intuitive. In contrast, the probability that there will be no volunteer is $(1 - p)^n$, and it follows from (3.4.1) that the Nash equilibrium probability of no volunteer is $(C/B)^{n/(n-1)}$, which is an *increasing* function of the number of potential volunteers, n . Moreover, $\lim_{n \rightarrow \infty} Pr\{\text{no volunteer}\} = C/B > 0$. Intuitively, one might expect that the greater the number of potential volunteers, the more likely it will be that *at least one* will actually volunteer.

Table 3.2. Frequencies of volunteer decisions and no-volunteer outcomes.
(Source: Frazen, 1995)

Group Size n :	2	3	5	7	9	21	51	101
Frequency of Volunteer Decisions:	.65	.58	.43	.25	.35	.30	.20	.35
Frequency of "No-Volunteer" Cases:	.12	.07	.06	.13	.02	.00	.00	.00

Table 3.2 shows the results of a one-shot volunteer's dilemma experiment, with $B = 100$ and $C = 50$ (Frazen, 1995). The volunteer rates in the middle row generally conform to the intuitive Nash predictions that an increase in n will decrease the chances than any single

person decides to volunteer. But the frequencies of a no-volunteer outcome, shown in the bottom row, decline to 0 as n is increased, which contradicts the Nash prediction that this frequency should be increasing and should converge to $C/B = 1/2$.

In a quantal response model, the injection of (enough) noise will reverse the unintuitive Nash prediction that the probability of not getting at least one volunteer increases in n :

Proposition 3.2. *In a quantal response equilibrium for the volunteer's dilemma game, the probability that no one will volunteer is decreasing in the number of potential volunteers for a sufficiently low precision parameter λ . Furthermore, for any finite λ , the limit of the probability of a no-volunteer outcome goes to 0 as $n \rightarrow \infty$.*

Proof (Goeree, and Holt, 2005). Recall that the expected payoff difference is $B(1-p)^{n-1} - C$, and therefore, equation (3.3.2) becomes: $F^{-1}(p)/\lambda = B(1-p)^{n-1} - C$. The first step is to express this QRE equilibrium condition in terms of the probability of getting no volunteer, denoted by $P = (1-p)^n$. Since the payoff shocks are i.i.d., the inverse distribution function is symmetric, i.e. $F^{-1}(p) = -F^{-1}(1-p)$. Now the equilibrium condition can be expressed:

$$-F^{-1}(P^{1/n})/\lambda = BP^{(n-1)/n} - C, \quad (4.2)$$

or equivalently,

$$P^{1/n} = F\left(\lambda[-BP^{(n-1)/n} + C]\right). \quad (4.3)$$

It is straightforward (but tedious) to use the implicit function theorem on (3.4.3) to establish that

$$\frac{dP}{dn} = -\frac{P \ln(P)}{n} \frac{\lambda B f\left(F^{-1}(P^{1/n})\right) - P^{-1+2/n}}{(n-1)\lambda B f\left(F^{-1}(P^{1/n})\right) + P^{-1+2/n}}. \quad (4.4)$$

To construct a proof by contradiction, note that dP/dn can only be non-negative if the numerator of the fraction on the far right side of (3.4.4) is non-negative, or if $1/\lambda \leq P^{1-2/n} B f\left(F^{-1}(P^{1/n})\right)$. The right side of this inequality is bounded by $B \max\{f\}$, so the inequality will be violated for sufficiently small values of λ , and $dP/dn < 0$ as a consequence. The final step is to show that P tends to 0 when n goes to infinity. Suppose, in contradiction, that $\lim_{n \rightarrow \infty} P > 0$. This implies that $P^{1/n}$ tends to 1, which contradicts (3.4.3) since the right side limits to $F\left(\lambda(-BP + C)\right) < 1$ for any finite λ . QED

The intuition behind this proposition is that, with noisy decisions, a rise in n only causes a slight reduction in probability of volunteering, which is more than compensated for by the fact that there are more potential volunteers who might be pushed in that direction by random elements.

The graphical methods introduced in this section and the previous one build on the Palfrey and Rosenthal (1985) model of voter participation with random shocks to the cost of voting. This general approach has been used to analyze a wide range of other binary choice games, including step-level public goods where the required number of “volunteers” is greater than one (Goeree and Holt, 2005) and voting participation games to be considered next.

5. Voting: Participation Games

This section will introduce the application of QRE methods to voting problems, where the payoff calculations must take into account the (typically small) effects of vote decisions on election outcomes. Of course, small differences can have large effects in a QRE model where choice probabilities based on small payoff differences are particularly sensitive to random shocks. The voting game to be considered is similar to that of Palfrey and Rosenthal (1985). In their model, there are two types of voters, with N voters of each type (so the number of players is $2N$). The outcome is determined by majority rule, with ties determined by the flip of a coin. There are two possible outcomes, with one outcome preferred by voters of one type, and the other outcome preferred by voters of the other type. The payoff to each voter is V if their preferred outcome is selected, and it is 0 otherwise. Each person who votes incurs a cost, c , irrespective of the outcome of the election.

In the absence of a cost of voting ($c = 0$), each person would always vote, and the outcome would be a tie in this symmetric model. We assume that the value of the preferred outcome is high enough so that it cannot be a Nash equilibrium for nobody to vote. To see what this implies, note that each side has a $1/2$ chance of winning if nobody votes, so the expected payoff from not voting is $V/2$. When none of the others are voting, one’s vote will swing the election, so the expected payoff will be $V - c$. Thus the expected payoff gain from voting when none of the others vote is $V/2 - c$, which is assumed to be positive to rule out the no-vote equilibrium. This assumption also implies that it is a Nash equilibrium for everyone to vote when there are equal numbers in each group. When everyone else is voting, not voting results in a payoff of 0, which is less than the expected payoff from voting to ensure a tie, $V/2 - c$.

There may also be symmetric mixed-strategy equilibria in which each person votes with a probability p , irrespective of their group identity. To analyze these equilibria, we need to derive formulas for the probabilities of specific vote sums, using binomial probability calculations. For example, with N voters of a type 2, the probability that there will be n_2 votes from this group is “ N choose n_2 ” times $p^{n_2}(1 - p)^{N - n_2}$, or $\binom{N}{n_2}p^{n_2}(1 - p)^{N - n_2}$. From the point of view of a voter of type 1, the relevant events also depend on the number of votes cast by the $n - 1$ other type 1 voters, since this number, along with n_2 , determines whether the voter’s decision to vote will either create a tie (raising the value of the outcome from 0 to $V/2$) or break a tie (raising the expected value from $V/2$ to V). Let the probability of creating a tie be denoted by P_T , and let the probability of breaking a tie

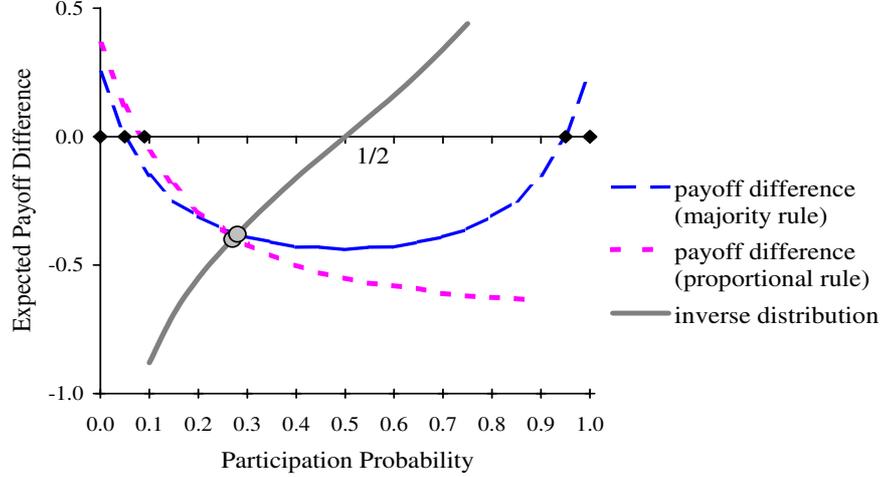


Figure 5: *Equilibrium points for a Symmetric Participation Game Under Two Alternative Voting Rules*
 (Key: Nash equilibria are diamonds, QRE are circles.)

be denoted by P_B , where both of these will be functions of N and p .⁶ It follows that the increase in expected payoff from voting is: $(P_T + P_B)V/2 - c$.

Figure 3.5 shows the expected payoff difference under majority rule as a function of the common participation probability, as shown by the “U-shaped” dashed line. This difference was calculated for the parameters used in an experiment reported by Schram and Sonnemans (1996b), with $c = 1$, $N = 6$, and $V = 2.5$ (Dutch guilders). The intuition behind this “U shape” is that the expected payoff difference is highest at each endpoint, i.e. when none of the others votes ($p = 0$), in which case a vote will result in a sure win, or when all of the others vote ($p = 1$) and a vote will create a tie. As before, the points where this payoff difference line is 0 correspond to Nash mixed equilibria; there is one with a low probability of voting, and one with a high probability (Palfrey and Rosenthal, 1985). This participation game has two interesting strategic dimensions. First, each person in a winning group has an incentive to “free ride” off of the others’ decisions to vote, just as players in a public goods game have an incentive to free ride on others’ contributions.⁷ Second, all members in

⁶A tie is created when there are n_2 votes on the other side and $n_2 - 1$ votes on one’s own side, so $P_T = \sum_{n_2=0}^N \binom{N}{n_2} \binom{N-1}{n_2-1} p^{2n_2-1} (1-p)^{2N-2n_2}$. A tie is broken when there are n_2 votes on the other side and n_2 votes on one’s own side, so $P_B = \sum_{n_2=0}^N \binom{N}{n_2} \binom{N-1}{n_2} p^{2n_2} (1-p)^{2N-2n_2-1}$.

⁷Public goods games will be discussed in the next chapter, where the focus is on the conflict between selfish motives to free ride and altruistic motives to contribute, a conflict that is studied in the lab by varying

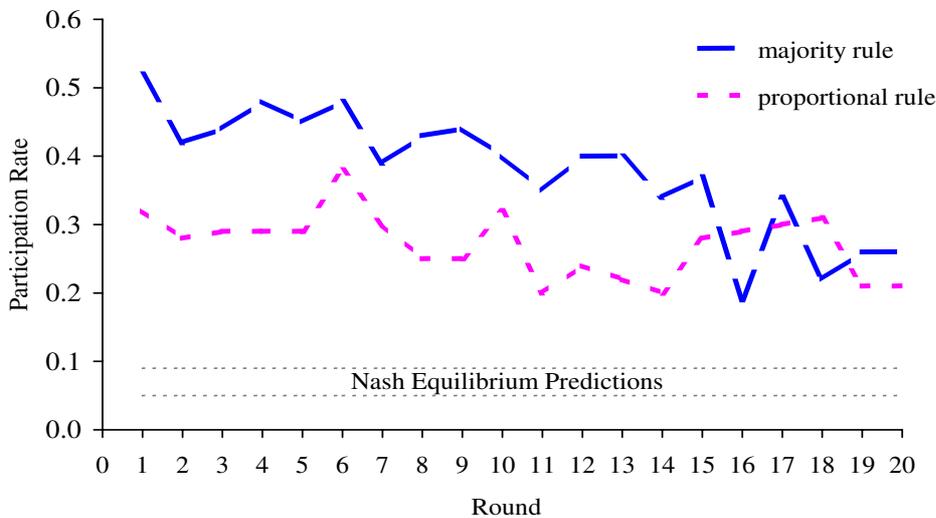


Figure 6: *Participation Rates Under Under Majority and Proportional Rule* (Source: Schram and Sonnemans 1996b)

a group might be better off if they can coordinate on a high-participation outcome for their group, although a symmetric high-participation equilibrium for both groups (the equilibria on the right side of Figure 3.5) is costly for both groups, and hence Pareto inferior.

Schram and Sonnemans used a design in which subjects were randomly rematched into competing groups of size 6 in a series of 20 voting participation games. The average participation rates under this majority rule treatment are shown in Figure 3.6 as the line with the longer dashes. Note that the participation rates start at about 0.5 and fall to the range from 0.2 to 0.3 in the final rounds. These rates are well above the Nash prediction of 0.05 for this treatment. To evaluate the QRE prediction, we must draw the inverse distribution function line in Figure 3.5. This line is plotted for a logit model with $\lambda = 0.4$, selected to results in an intersection with the expected payoff difference line at a participation rate in the observed range, between 0.2 and 0.3.⁸

Schram and Sonnemans implemented a second treatment in which each group's share of a prize, V , was equal to its share of the total votes cast. As before, voting entails a cost c that is the same for all, and there are N voters in each group. Consider the decision of a

the incentives and the estimating parameters that determine the degree of altruism.

⁸Even though the subjects in the experiments all faced the same, deterministic cost of voting c , the theoretical analysis of the quantal response equilibrium would coincide with a Bayesian Nash equilibrium in which each voter faces a randomly determined, i.i.d. voting cost with mean c . In other words, the epsilon shocks in the QRE model could be interpreted as cost shocks, which is the approach taken in Palfrey and Rosenthal (1985).

voter in group 1. If n_1 votes are cast by the other members of group 1 and n_2 votes are cast by those in group 2, then the change in group 1's vote share caused by an additional vote is $\frac{n_1+1}{n_1+n_2+1} - \frac{n_1}{n_1+n_2}$. The expected change in vote share, denoted $\Delta(N, p)$, is calculated in a standard manner.⁹ With this notation, the increase in expected payoff from voting can be expressed as a function of N and p : $\Delta(N, p)V - c$. The line with short dashes in Figure 3.5 shows this expected payoff difference for the parameters used in this treatment ($N = 6$, $V = 2.22$ and $c = 0.7$). The line is a smoothly decreasing function of p because each vote has less effect on the proportion when the others are more likely to vote. The Nash equilibrium for these parameters is 0.09, as indicated by the diamond in the figure where the expected payoff difference crosses the 0 axis, and the quantal response equilibrium for $\lambda = 0.4$ is at essentially the same point as was the case for the majority voting rule treatment. The data from the experiment are consistent with these predictions. Under proportional voting, the participation rates (shown by the short dashed line in Figure 3.6) start lower than was the case under majority voting, but the rates for the two treatments end up being approximately the same, in the 0.2-0.3 range.

The declining participation rates observed in both treatments could be due to a dynamic process in which subjects are responding to the relatively low impact of their vote decision on the payoff outcome. For the middle range of participation rates (in the center of Figure 3.5), the expected payoff difference lies below the inverse distribution line, so the quantal responses to these low incentives to vote would be to reduce participation rates. Thus the quantal response analysis explains both the falling participation rates and the tendency for these rates to stay well above the Nash predictions.

This analysis can be generalized to the case of two groups of differing sizes, so that any equilibrium may involve differing participation rates, p_1 and p_2 , which are determined by solving simultaneous equations. Cason and Mui (2003) used the asymmetric participation game model to analyze strategic interactions between majority and minority groups in a laboratory experiment. In the "certain roles" treatment, all subjects knew their own roles with certainty. A majority (3 of 5 voters) were given payoffs that caused them to favor the "reform" option over the "status quo," and the minority had opposing preferences. All voters were given the same voting cost, regardless of group. Subjects were randomly rematched in groups of size 5 in a series of rounds, with the reform decision being determined by a majority-rule vote in each round. The cost of voting was varied across treatments to test the Nash equilibrium prediction that participation rates for both rates would be inversely related to the cost of voting. This negative relationship between voting costs and participation rates was observed for the majority group, but not for the minority group. Cason and Mui use numerical calculations to show that these conflicting patterns are both consistent with a quantal response equilibrium with a moderate amount of noise. The quantal response equilibrium also predicted the directions of deviations from Nash participation rate predictions (extreme predictions were "pulled" in the direction of a fifty-fifty split). Finally,

⁹Using the formula for a binomial distribution, the expected value of the increase in one's vote share is: $\Delta(N, p) = \sum_{n_2=0}^N \sum_{n_1=0}^{N-1} \binom{N}{n_2} \binom{N-1}{n_1} \left[\frac{n_1+1}{n_1+N_2+1} - \frac{n_1}{n_1+n_2} \right] p^{n_1+n_2} (1-p)^{2N-n_1-n_2-1}$, where the outside sum is taken with respect to the number of votes from the other group and the inner sum is taken with respect to the number of votes in one's own group (1).

any incidence “incorrect” votes for the less preferred alternative cannot be explained by a theory that does not incorporate random elements.

The main focus of the Cason and Mui paper was on a comparison of the “certain roles” treatment (just discussed) with one in which some voters faced uncertainty about whether or not they would benefit from the reform. Specifically, two voters knew that they preferred the reform, but the other three only knew that one of them (selected at random *ex post*) would benefit from the reform. The payoff parameters were selected so that a majority would always favor reform *ex post*, but the uncertainty would cause a majority to vote against the reform if voting were costless. With costly voting, the Nash (and QRE) equilibria involve mixed strategies, and the predicted probabilities of reform were lower with role uncertainty. This reduction in the incidence of reform under uncertainty was observed in the data.

The analysis of asymmetric participation games suggests a wider range of applications. For example, political scientists are interested in studying other factors that may increase the chances that a revolution or other radical reform may occur, but such events are often decided in contest settings that do not correspond to majority voting. It may be reasonable to model the probability of a successful revolution as an increasing function of the number who incur a cost to support the revolution, but this function may not have a discrete jump in value from 0 to 1 as that number switches from a minority to a majority. In particular, a proportional voting rule would implement the assumption that the chance of success is equal to the proportion of support among activists (those who incur the cost of supporting their preferred alternative).¹⁰ Cost differences or asymmetries in the probability-of-success function could be introduced to induce a bias in favor of (or against) the *status quo*. In this case, the revolution could be modelled as an asymmetric participation game with a modified proportional voting rule. As noted above, the interesting strategic features of this game are that it incorporates elements of both public goods and coordination games, since each person would prefer to “free ride” off of the costly activities of others in their group, but all people in a group might be better off if they could coordinate on high activity levels.

6. Jury Voting

Elections are typically thought of as mechanisms for settling differences among individuals who do not agree on particular issues. There is, however, a second function of voting procedures: the aggregation of information across individuals who may agree on policy matters but who have different sources of information. The idea that elections aggregate information dates to the eighteenth century writings of Condorcet, and has generated a large body of research since then. The main insight of this literature can be illustrated with a simple example. Suppose that N members of a jury observe a trial and form independent opinions about the guilt of the defendant. If the defendant is actually guilty and the trial is informative, then it is more likely that jurors will vote to convict. The way the trial

¹⁰This proportional success probability function is widely used in the rent-seeking and contest literatures, which will be discussed in the final section of this chapter.

information is interpreted is noisy, so some jurors may vote either way, but there will be a smaller tendency to vote for a conviction if the defendant is actually innocent. Let $P(V_C|G)$ be the probability that a juror will vote to convict, V_C , if the defendant is guilty, and let $P(V_C|I)$ be the probability that a juror will vote to convict if the defendant is innocent. In a large jury with independent decisions, the fraction who vote to convict will be approximately $P(V_C|G)$ if the defendant is guilty, and this fraction will be approximately $P(V_C|I)$ otherwise, with $P(V_C|I) \leq P(V_C|G)$. Now let ρ be the fraction of the jury that is required to vote for conviction in order to obtain a conviction, e.g. $\rho = 0.5$ under majority rule and $\rho = 1$ under unanimity. If ρ is set too low, below $P(V_C|I)$, then innocent individuals would tend to get convicted. Obviously, one would like to set the required fraction to be between $P(V_C|I)$ and $P(V_C|G)$, but these cutoffs may not be known with precision and may differ from trial to trial. In any case, the chances of convicting an innocent person can be reduced by using a higher value of the cutoff fraction ρ , *if these probabilities are not affected by the voting rule cutoff itself.*

This intuition, that a higher standard of conviction will tend to protect the innocent, is implicitly based on a naive model of voting, i.e. that people form opinions based on the trial evidence and vote without any strategic consideration for how others might vote. Naive voting, however, may not be a rational way to use one's information if unanimity is required (Austin-Smith and Banks, 1996; Feddersen and Pessendorfer, 1996). To see this, suppose that each person receives a signal, g (guilty) or i (innocent), with a guilty signal being more likely if the defendant is guilty, and an innocent signal being more likely if the defendant is innocent. If each juror votes naively (to convict if they see g and to acquit if they see i), then consider the dilemma of a juror who sees an i signal. Under unanimity, a vote to acquit will preclude conviction, regardless of how the others vote. Even a vote to convict will not result in conviction if at least one of the others voted to acquit. The only way that a person's vote will matter, then, is if *all* of the others voted to convict.¹¹ Under naive voting, this means that all of the $n - 1$ others saw a g signal, so a (naive!) juror who really believes that the others are voting naively, might want to vote to convict, contrary to the signal. In this case, naive voting is not a best response to naive voting, so voting naively would not be a Nash equilibrium.¹² Moreover, the incentive for a juror with an i signal to deviate from naive behavior by voting to convict is stronger for a larger jury. In the Nash equilibrium to be explained next, the probability of voting to convict given an i signal will be an increasing function of n , the jury size.

Rational Strategic Voting

¹¹In many contexts, there may be three possible outcomes: conviction, acquittal, and mistrial, which may occur if there is not total agreement among the jurors. Following Feddersen and Pessendorfer (1998), we will not consider the possibility of mistrials.

¹²The key insight here is that being a "swing voter" is an informative event, and a voter should take this into account in making a decision, just as a bidder in a first-price, common-value auction should realize that a bid will only be relevant if it is the high bid, which may mean that the person over-estimated the value of the prize. Failure to adjust for this implicit information may cause the winning bidder to lose money due to the "winner's curse." Similarly, failure to condition the voting decision under unanimity properly may lead to a "swing-voter's curse" (Feddersen and Pessendorfer, 1996).

Consider a game in which the defendant is guilty (state G) with probability $1/2$ and innocent (I) with probability $1/2$.¹³ Each of the n jurors receive an independent Bernoulli signal that is g with probability p if the defendant is guilty, and is i with probability p if the defendant is innocent, where $p > 1/2$. After observing their own private signals, jurors vote to convict or acquit. The outcome under unanimity is a conviction, C , only if all n jurors vote V_C . Jurors have identical utilities that depend on the state and outcome, and on a parameter q that determines the relative cost of making a mistake in either direction; for a false conviction, $u(C, I) = -q$, and for a false acquittal, $u(A, G) = -(1 - q)$. The utilities for the correct decision are normalized to 0: $u(A, I) = u(C, G) = 0$. Let $\hat{P}(G)$ be the voter's posterior probability of G , given relevant information (as will be explained below). Then the expected payoffs for a jury decision to convict or acquit, denoted by π_C and π_A respectively are: $\pi_C = 0\hat{P} - q(1 - \hat{P})$ and $\pi_A = -(1 - q)\hat{P} + 0(1 - \hat{P})$, so:

$$\pi_C - \pi_A = \hat{P}(G) - q. \quad (6.1)$$

Thus the preferred outcome is conviction when the posterior probability of guilt is above q , which represents a “threshold of reasonable doubt.” In what follows, we assume that $p > q$, which implies a jury of one would vote to convict with a guilty signal and to acquit otherwise.

In this game, a strategy is a probability of voting for conviction, conditional on the private information: $\sigma : \{g, i\} \rightarrow [0, 1]$. There are two ways that a vote for conviction can be generated, depending on whether the signal is “correct” or not. Given the probability, p , of a correct signal, the conditional probabilities of a vote to convict are:

$$P(V_C|G) = p\sigma(g) + (1 - p)\sigma(i) \quad \text{and} \quad P(V_C|I) = (1 - p)\sigma(g) + p\sigma(i). \quad (6.2)$$

The symmetric equilibrium to be constructed will have the property that those with g signals always vote to convict, $\sigma(g) = 1$, and those with i signals vote to convict with a probability: $\sigma(i) < 1$, so the equations in (3.6.2) can be expressed:

$$P(V_C|G) = p + (1 - p)\sigma(i) \quad \text{and} \quad P(V_C|I) = (1 - p) + p\sigma(i). \quad (6.3)$$

Next consider the decision for a juror who receives the i signal. Under unanimity voting, the only case in which a vote will matter is when the $n - 1$ others vote to convict, so the juror should evaluate the probability of guilt conditional on both an i signal and on $n - 1$ other V_c votes. Given the simplifying assumption that each state (G or I) is equally likely *a priori*, the $1/2$ terms cancel out of the numerator and denominator of Bayes' rule, and the conditional probability of G is:

¹³The analysis that follows can be generalized to allow for unequal prior probabilities, but we will use equal priors to convey the main results with a minimum of notation.

$$\widehat{P}(G|i, n-1 \text{ other } V_C \text{ votes}) = \frac{(1-p)P(V_C|G)^{n-1}}{(1-p)P(V_C|G)^{n-1} + pP(V_C|I)^{n-1}} = q, \quad (6.4)$$

where the $(1-p)$ terms represent the probability of getting an i signal when the defendant is guilty, and the p term in the denominator is the probability of i signal when the defendant is innocent. The final equality in (3.6.4) follows from (3.6.1) and the requirement that the voter must be indifferent between the two outcomes in a mixed-strategy equilibrium. (More precisely, indifference between the two outcomes implies indifference between the two vote decisions that might affect these outcomes.) Equation (3.6.4) can be solved for the probability of a vote to convict, conditional on the defendant being innocent:

$$P(V_C|I) = R_n P(V_C|G), \quad (6.5)$$

where the ratio of conviction vote probabilities is:

$$R_n = \left(\frac{(1-q)(1-p)}{pq} \right)^{\frac{1}{n-1}}. \quad (6.6)$$

Then (3.6.5) can be used with (3.6.3) to solve for the equilibrium probability:

$$\sigma(i) = \frac{R_n p - (1-p)}{p - R_n (1-p)}, \quad (6.7)$$

where R_n is expressed in terms of exogenous parameters in (3.6.6).^{14, 15}

Consider what happens to the probability of convicting an innocent defendant as the jury size grows. As n tends to infinity, R_n tends to 1 and $\sigma(i)$ tends to 1 in (3.6.7). In this case, it follows from (3.6.2) that the probabilities of a vote to convict in either state converge to 1. These tendencies suggest that unanimity rules may not protect innocent defendants as the jury size increases. More precisely, the probability of convicting an innocent defendant is the

¹⁴By construction, a juror with an i signal is indifferent and can do no better than to randomize. The final step is to show that a juror with a g signal would vote to convict with probability 1 if jurors with the i signal are randomizing. Suppose (in contradiction) that those with g signals were to randomize in equilibrium, so they must be indifferent between voting to convict or acquit. But then those with i signals would not be indifferent; they would always vote to acquit, so the only votes to convict would be cast by those with g signals. Then a person with a i signal would only be a swing voter if all others have g signals, and hence this person would prefer to vote to convict, which is a contradiction.

¹⁵As is the case with many voting models, there may be “trivial” Nash equilibria that are due to the fact that a single voter may not be able to affect the outcome. In this case, one of these equilibria is for the strategy as a function of the signal to vote convict with probability 0: $\sigma(i) = 0$, $\sigma(g) = 0$. As long as there are at least 2 voters, the outcome will always be an acquittal, and unilateral deviations by a single voter cannot change this. These Nash equilibria are unintuitive, and hence attention is restricted to symmetric, informative equilibria in which the probability of a vote to convict is higher with a g signal than with an i signal: $\sigma(g) > \sigma(i)$.

probability that all N jurors vote to convict, $P(C|I)^n$, and this probability can be calculated by substituting the solution for $\sigma(i)$ into the right equation in (3.6.3) and raising the result to the power n . If $q = 0.5$ and $p = 0.7$, for example, the Nash equilibrium probability of convicting the innocent under unanimity is 0.14 with a jury of size 3, and it *increases* to 0.19 with a jury of size 6. Moreover, Feddersen and Pessendorfer (1998) show that the limiting value of this error probability is 0.23 as N goes to infinity for this example. In fact, they show that this limit is strictly positive for all parameter values under consideration, i.e. for $1 - p < q < p$.¹⁶

In contrast, consider what happens under majority rule as the jury size increases. If mistakes in either direction are equally costly ($q = 1/2$), then there is an equilibrium in which jurors vote naively, with $\sigma(g) = 1$ and $\sigma(i) = 0$. The intuition is that the swing-voter's curse has no "bite" in this case, since being pivotal means that the others are more-or-less evenly split, and hence this event is not informative.¹⁷ Thus a "naive" vote would result from the appropriate use of private information. With naive (and in this case rational) voting under majority rule, the intuition for the Condorcet information aggregation result holds, and large juries would tend to make no mistakes. From a game-theoretic perspective, what was wrong with the original intuition drawn from the Condorcet jury voting literature was the assumption that the nature of voting behavior is unaffected by the voting rule itself. Whether actual behavior is naive or not in specific settings, however, is a question that can be addressed with laboratory experiments, which is the next topic.

Experiment Results

Guarnachelli, McKelvey, and Palfrey (2000) report a laboratory experiment done with neutral terminology, $p = 0.7$, $q = 0.5$, and juries of size 3 and 6. Here we focus on the results for the unanimity rule (they also ran some trials with a majority rule). Subjects were recruited in groups of 12 and were randomly matched into juries in a sequence of jury meetings, with payoffs that were contingent on the jury decision and the state, which was announced after each round. With these parameters, the Nash equilibrium mistake probabilities are shown in Table 3.3. The prediction is that innocent defendants will not be protected under unanimity for juries of size 6 and above.

Table 3.3. Mistake Probabilities in a Nash Equilibrium ($q = 0.5, p = 0.7$)

	$n = 3$	$n = 6$	$n = 12$	$n = 24$
$P(C I)$ under unanimity	.14	.19	.21	.22
$P(C I)$ under majority	.22	.07	.04	.01
$P(A G)$ under unanimity	.50	.48	.48	.47
$P(A G)$ under majority	.22	.26	.12	.03

One striking result of the experimental data is that subjects tend to vote strategically

¹⁶The formula for this limit is $((1 - q)(1 - p)/qp)^{\frac{p}{2p-1}}$.

¹⁷Depending on whether the number of jurors is even or odd, being pivotal might mean breaking or creating a tie.

(not naively) when the Nash prediction is that they will vote strategically. Under unanimity, the fraction of votes to convict after receiving an i signal goes from 0.36 for juries of size 3 to 0.48 for juries of size 6. These proportions are roughly consistent with the corresponding Nash predictions of 0.31 and 0.65 respectively, which are derived from equation (3.6.7). Thus the data are qualitatively consistent with the game-theoretic intuition that jurors with innocent signals will be more likely to vote to convict in larger juries, where the swing-voter's curse has more of an effect. The results contrast sharply with the naive voting rule, where the probability of voting to convict is 0 for all jury sizes.

Next consider the unintuitive Nash prediction, that an increase in the size of the jury will *increase* the probability of conviction under unanimity. This did not happen. In fact, the proportion of incorrect convictions of innocent defendants went from 0.19 down to 0.03, instead of rising as predicted in a Nash equilibrium. See the first two columns of Table 3.4 for a comparison of the data and Nash predictions for the rate of incorrect convictions under unanimity voting.

Table 3.4 Proportions of incorrect convictions under unanimity.
(Source: Guarnachelli, McKelvey, and Palfrey, 2000)

	Data	Nash	QRE
Smaller Jury ($n = 3$):	0.19	0.14	0.19
Larger Jury ($n = 6$):	0.03	0.19	0.07

The Equilibrium Effects of Noise in Jury Voting

For a given private signal, voting is a binary decision, so we use the approach taken in the previous two sections in which the probability of a vote to convict is an increasing function of the expected payoff difference for the two decisions, conditional on the private signal, $s \in \{i, g\}$, and on the strategies used by the $n - 1$ others. Let these two expected payoffs be denoted by $\pi_{V_C}(s, n - 1, \sigma(i), \sigma(g))$ and $\pi_{V_A}(s, n - 1, \sigma(i), \sigma(g))$. Then the equations analogous to (3.5.2) are:

$$F^{-1}(\sigma(i))/\lambda = \pi_{V_C}(i, n - 1, \sigma(i), \sigma(g)) - \pi_{V_A}(i, n - 1, \sigma(i), \sigma(g)), \quad (6.8)$$

$$F^{-1}(\sigma(g))/\lambda = \pi_{V_C}(g, n - 1, \sigma(i), \sigma(g)) - \pi_{V_A}(g, n - 1, \sigma(i), \sigma(g)). \quad (6.9)$$

Solving these equations requires calculation of the probability that a change in one's vote from V_A to V_C will alter the outcome from A to C . Then the expected utilities for these outcomes can be used to calculate the expected payoff differences on the right sides of (3.6.8) and (3.6.9). Even without detailed calculations, we can obtain a useful limit result, which is essentially the same as Theorem 1 in Guarschelli, McKelvey, and Palfrey (2000), who used the logit form of the response function.

Proposition 3.3. Fix $\lambda < \infty$. Under unanimity, for every δ , there exists an number $N(\delta, \lambda)$ such that for all $n > N(\delta, \lambda)$, the probability of conviction in any quantal response equilibrium is less than δ regardless of whether the defendant is innocent or guilty.

Proof. Since all outcome-contingent payoffs are less than one, it must be the case that $\pi_{V_C}(i, n-1, \sigma(i), \sigma(g)) - \pi_{V_A}(i, n-1, \sigma(i), \sigma(g)) < 1$. Then the equilibrium conditions, (3.6.8) and (3.6.9), imply that $F^{-1}(\sigma(i)) > \lambda$ and $F^{-1}(\sigma(g)) > \lambda$. It follows that both $\sigma(i)$ and $\sigma(g)$ are greater than $F(\lambda)$, which is strictly positive when the preference shocks have positive density on the real line. (In a logit equilibrium, this bound is $1/(1+e^\lambda)$.) Since this lower bound on the acquittal probabilities is independent of n it follows that the probability of obtaining at least one vote to acquit tends to 1 as n tends to infinity. QED

To summarize, the Nash equilibrium conviction probability for an innocent defendant under unanimity is bounded away from zero for any jury size, which is inconsistent with the tendency for the incidence of false convictions to decrease as the jury size is increased in the experiment. In contrast, the quantal response prediction is that the probability of conviction goes to zero under unanimity as the jury size increases, due to noise effects. The QRE predictions for the parameters and estimated precisions from the Guarnachelli, McKelvey, and Palfrey experiment, shown in the far right column of Table 3.4, exhibit this pattern, which is consistent with the trend observed in the data column of the table.

7. The Traveler's Dilemma

Basu (1994) first described an interesting social dilemma in which the logic of standard game theory with perfectly rational players seems to be implausible. The game is motivated by a story of two travelers who have purchased identical antiques while on a tropical vacation. When their luggage is lost on the return trip, the airline claims representative informs them:

We know that the bags have identical contents, and we will entertain any claim between \$2 and \$100, but you will each be reimbursed at an amount that equals the *minimum* of the two claims submitted. If the two claims differ, we will also pay a reward of \$2 to the person making the smaller claim, and we will deduct a penalty of \$2 from the reimbursement to the person making the larger claim.

Each traveler would presumably know the cost of the lost antique, which would provide a natural focal point. But if one expects the other to claim this cost (or to claim \$100 if the cost was above the upper limit), then that person would have an incentive to undercut the anticipated claim for the other person. In fact, each person has an incentive to undercut any common claim that is above \$2, so the unique Nash equilibrium for this game is for each to claim the minimum amount, which may seem implausible given the relatively small penalty of \$2 for offering the higher claim.

The Traveler's Dilemma is richer than a Prisoner's Dilemma in the sense that the Nash decision is not a dominant strategy; the best claim to make would be just below the other's claim if it were known. In particular, if the other person is expected to claim \$100, then the best response is to claim \$99 (if claims are restricted to integer amounts). Obviously, there is no belief about the other's claim that would justify making a claim of \$100 in this game. If each player realizes that the other is thinking in this manner, then each will not expect the other to make a claim above \$99, but the best response to this upper limit on anticipated claims would then be \$98. Reasoning iteratively, one can rule out all claims above the minimum level on the assumption that it is common knowledge that each player is perfectly rational. As Basu (1994) notes, the Nash equilibrium is the unique *rationalizable* equilibrium in this game (with discrete decisions). These iterated-rationality arguments have some appeal, but note that the same arguments would apply if the penalty/reward amount were raised from \$2 to \$20 or lowered to \$1.01. Basu conjectured that claims would not fall to Nash levels for low values of the penalty-reward parameter, but he did not provide a formal theoretical explanation that was sensitive to these payoff parameter changes. A different approach might be to consider a model of adjustment, e.g. Cournot best responses to the most recently observed decision of the other player. But the Cournot best response to the other's claim is always to offer a slightly lower claim, regardless of the size of the payoff parameter, so this approach will not provide intuitive predictions, at least not without some modifications to be discussed below.

Experimental Data

The invariance of the Nash equilibrium to variations in the penalty/reward parameter is the motivation behind the laboratory experiment reported in Capra, *et al.* (1999). Subjects were recruited from economics classes at the University of Virginia. There were six sessions, with the number of participants in each session ranging from 9 to 12. Payoffs were scaled down so that earnings would be appropriate for student subjects. In particular, the instructions stipulated that people would be randomly paired and would have to make a claim decision on the interval [\$0.80, \$2.00]. Each session was characterized by a single penalty/reward parameter for 10 rounds of play, with random re-matching of participants after each round. The penalty/reward parameters used for the six sessions were: \$0.05, \$0.10, \$0.20, \$0.25, \$0.50, and \$0.80. As indicated above, the Nash equilibrium for each of these cases is the minimum claim of \$0.80.

The salient feature of the data for this experiment is that the observed claim levels responded sharply to changes in the penalty/reward parameter. The solid lines at the bottom of Figure 3.7 show the average claims by round for the two sessions with the highest R values (50 and 80 cents), and the two dashed lines at the top of the figure show the data for for the lowest R values (5 and 10 cents). With high values, the average claims started above \$1.20 in the initial round, but fell to levels very near the Nash prediction of \$0.80 in the final round. In contrast, the average claim data for sessions with relatively low parameter values started at about \$1.80 in the initial round, well away from the Nash prediction, and these averages leveled off at levels even farther from this prediction. Two additional sessions, with penalty-reward values of \$0.20 and \$0.25, produced average claims in an intermediate range

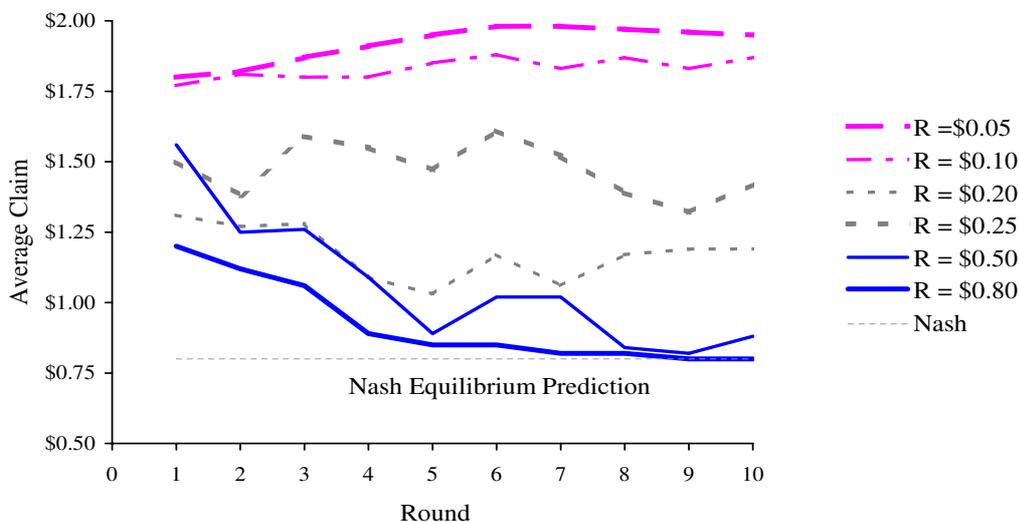


Figure 7: Average claim data for the traveler’s dilemma experiment. (Source: Capra *et al.* 1999)

between those with high values and those with low values, as can be seen from the dotted lines in the figure. Note that there is one “inversion” here in the sense that the average claims for the \$0.20 session were lower than the average for the \$0.25 session.

This experiment was somewhat unusual in that a whole range of payoff parameters was used, instead of restricting attention to a couple of parameter values, each with multiple sessions. Nevertheless, it is possible to construct a nonparametric test of the null hypothesis that changes in the payoff parameter have no effect, as predicted in a Nash equilibrium. The alternative hypothesis is that decreases in the payoff parameter will increase average claims. The most extreme observation would be for the average claims in the second row of Table 3.5 to increase from left to right, but there is one reversal.

Table 3.5 Average claims and theoretical predictions for the traveler’s dilemma. (Capra *et al.*, 1999)

Penalty/Reward Parameter:	\$0.80	\$0.50	\$0.25	\$0.20	\$0.10	\$0.05
Average Claim (last 5 rounds):	\$0.82	\$0.92	\$1.46	\$1.16	\$1.86	\$1.96
Logit QRE Prediction for $\lambda = 1.2$):	\$0.88	\$0.95	\$1.33	\$1.49	\$1.74	\$1.83
Nash Prediction:	\$0.80	\$0.80	\$0.80	\$0.80	\$0.80	\$0.80

There are $6! = 720$ different ways that the average claims could have been ranked, and of these, only 6 are either as extreme (one reversal in adjacent columns) or more extreme (no

reversals). The chances of observing a pattern at least as extreme as the one reported are 6 in 720, and therefore, the null hypothesis can be rejected at the one percent level. To summarize, the Nash equilibrium fails to explain the most salient feature of the data, i.e. the responsiveness of average claims to changes in the penalty/reward parameter.

Equilibrium Claim Distributions

The claim averages for the sessions shown in Figure 3.5 tend to level off after about 5 rounds, which suggests that an equilibrium has been approximately reached. In order to explain the logit QRE equilibrium predictions, it is necessary to introduce some notation. The Traveler’s Dilemma with a penalty/reward parameter of R is a symmetric, two-person game, and hence, we will consider a symmetric equilibrium in which each player uses the same mixed strategy that associates a probability σ_j with decision j , $j = 1, 2, \dots, J$, where J is the number of possible claim decisions. For example, the claims could be restricted to integer numbers of pennies on the interval $[80, 200]$. (Equilibria with continuous distributions of claims will be considered subsequently.) For simplicity, we will denote the claim amount with index j by x_j , e.g. $x_1 = 80$, $x_2 = 81$, etc. First, consider the payoff function for a player who selects claim x_j . The actual payoff will depend on the other player’s claim, x_k . This payoff, denoted by $v(x_j, x_k)$, will be $x_j + R$ if $x_j < x_k$ (reward obtained), it will be the claim of x_j if $x_j = x_k$ (no penalty or reward), and it will be $x_j - R$ if $x_j > x_k$ (penalty paid). Let p_k denote the probability representing a player’s beliefs that the other will choose decision k . The expected payoff for selecting decision j , denoted π_j , is then computed in a straightforward manner, and these payoffs are used to specify the logit response functions in equation (3.1.2), which are solved simultaneously to determine the logit equilibrium probabilities.

Figure 3.8 shows the average claim predictions for two of the treatments, one with a low R of \$0.10, and another with a high R of \$0.50, for a range of λ values. When the value of the precision parameter is close to zero (on the left side of the figure), the high noise causes the equilibrium distributions to “flatten out” to a uniform distribution on the whole range of claims, which in turn causes logit average claim predictions for all treatments to go to the midpoint of the interval, \$1.40. As the precision parameter is increased (moving to the right in the figure), the predictions diverge, with high claims predicted for the $R = \$0.10$ treatment and low claims predicted for $R = \$0.50$ treatment. For very high values of λ on the right side of the figure, the predictions for both treatments converge to the Nash equilibrium of \$0.80. Thus the qualitative features of the experimental data (the spread) are accounted for by the logit equilibrium for a wide range of precision values.

To summarize, the logit equilibrium probabilities depend on the precision parameter λ and on the properties of the payoff function. These predictions are more credible if the probabilities associated with decisions that are observed in the experiment are not relatively small. In the next chapter, we will discuss how the precision parameter can be estimated to maximize the likelihood of the observed data from the experiments. For present purposes, it is sufficient to note that the maximum likelihood estimate of λ was 1.2 for the data reported in Capra *et al.* (1999), with payoffs measured in pennies and using data for the final five rounds. This particular parameter value is then used to calculate the equilibrium

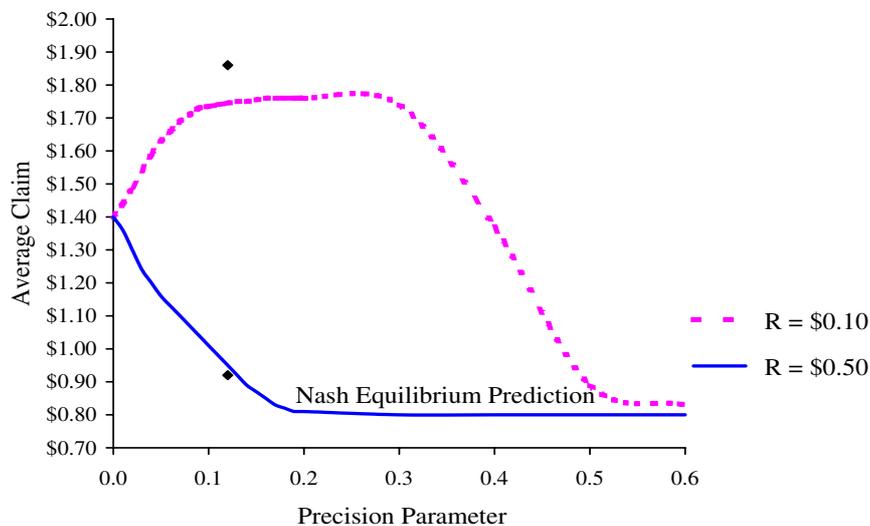


Figure 8: *Logit correspondence for selected high ($R = \$0.50$) and low ($R = \0.10) penalty-reward parameters in a traveler's dilemma. (Key: Diamonds show the data averages.)*

probabilities associated with each decision, which in turn can be used to generate a prediction of the average claim. These predictions are shown in the third row of Table 3.5, just above the Nash prediction, which is \$0.80 for all treatments. The QRE predictions correctly indicate that claims will converge to near-Nash levels for relatively high values of R and will cluster at the *opposite* end of the set of feasible claims for relatively low values of R . The diamond points in Figure 3.8 show the logit prediction parameter estimate (horizontal coordinate) and the data averages in rounds 5-10 for the two treatments (vertical coordinate). These data points lie fairly close to the logit prediction lines in the figure.

It is instructive to think about the equilibrium claim distributions as the results of an iterated process of leaning and noisy response. In Figure 3.7, with a high penalty/reward parameter of 50 cents, for example, the average claims start near the middle of the range at about \$1.20, but as subjects learn to lower claims to get a better chance for the 50 cent reward, the average claims fall in successive rounds. In this sense, the empirical distribution of claims would tend to fall, becoming more concentrated near the Nash equilibrium. Conversely, the initial claims are higher for low values of the R parameter, and the small penalty for being low does not induce a fall in claims. In fact, the average claims tend to move away from the Nash prediction in the direction of the maximum claim for several rounds, before levelling off.

An admittedly crude model of the learning process at the aggregate level could be that subjects initially have relatively dispersed beliefs, with the p_j probabilities being of uniform

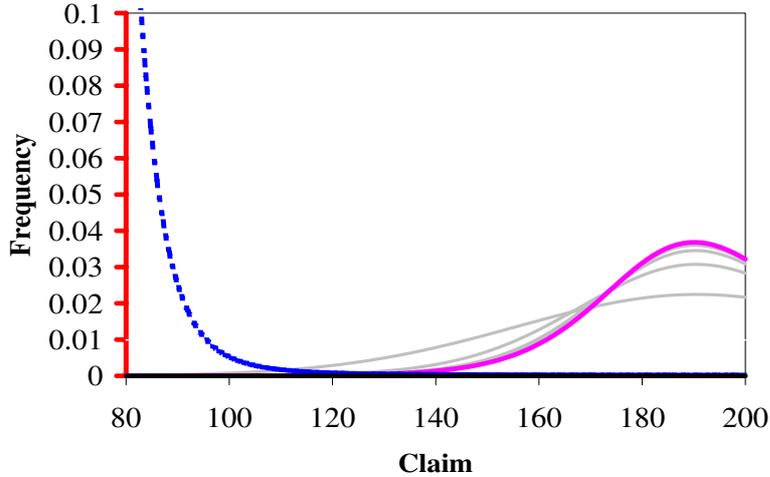


Figure 9: *Logit QRE claim distributions for the traveler's dilemma with $\lambda = 0.1$, for $R = 50$ (dashed line) and $R = 5$ (dark line).*

size for the whole range. With uniform beliefs and a low R of 5 cents, the noisy best response determined by equation (3.1.2) (with $\lambda = 0.1$) is no longer uniform, but is a flat hill with a peak in the range of high claims, as shown by the flatter of the light gray lines in Figure 3.9. Then if this line is used to generate the belief distribution, i.e. the p_j probabilities in (3.1.1), then the noisy response probabilities determined by equation (3.1.2) yield the gray line that is somewhat more peaked. After three iterations, all of the successive gray lines overlap and are essentially at the equilibrium distribution shown by the dark curve in the figure.¹⁸ Although this iterated, aggregate process of learning and noisy responses is instructive, several more detailed models of learning based on individual experiences will be developed in a later section.

The Capra *et al.* (1999) experiment involved random matching, a protocol that is intended to allow learning from experience but to reduce opportunities for collusion, which would be more likely if individuals were matched with the same person in all rounds. We conjecture that claims would be uniformly higher for most treatments under fixed matchings. In the other extreme case, consider a situation in which each person only plays a single round of the Traveler's Dilemma game. Data for one-shot games of this nature, reported in Goeree and Holt (2001), confirm the observation that the Nash prediction is accurate for high values of the penalty-reward parameter, and that claim data are clustered at the upper end of the claim distribution, far from the Nash prediction, for games played with low values of this parameter. It is not reasonable to expect the equilibrium convergence of beliefs and actions

¹⁸See Holt (2005, Chapter 11) for the details of setting up a spreadsheet program that lets one experiment with these iterations for various initial belief distributions and parameter values.

in a single shot game, as required by equation (3.1.3), but the iterated sequence of gray noisy response lines in Figure 3.9 indicates how even one noisy response to diffuse prior beliefs might generate data far from the Nash prediction in a single-shot game. The conjecture that initial prior beliefs are uniformly distributed over the range of claims, independently of the structure of the game (the R parameter) is obviously simplistic. This raises an interesting question of how beliefs are formed on the basis of *introspection* alone, in the absence of any observations of others' decisions. Models of introspection could be constructed on the basis of iterated noisy thinking about what the other person will do, what they think I will do, what they think I think they will do, etc. Such models will be considered in a later chapter.

Comparative Static Effects

The calculated effects of changing the penalty/reward parameter, as shown in Figures 3.8 and 3.9, were done for specific parameter values, which raises the question of general comparative statics effects of changes in R . Since the claim choices in the experiment were not restricted to integer values, we will consider a model in which the claim, x , can be any real number on the interval $[\underline{x}, \bar{x}]$. Let the equilibrium distribution of x be denoted by the density $f(x)$, with distribution function $F(x)$. Thus the expected payoff is:

$$\pi(x) = \int_{\underline{x}}^x (y - R)f(y) dy + (x + R)[1 - F(x)], \quad (7.1)$$

where the first term on the right side of (3.7.1) corresponds to the case where the player's claim is high, so the minimum is determined by the other's claim and a penalty is subtracted. The second term on the right side of (3.7.1) corresponds to the other case in which the player's claim is low, so it determines the minimum and a reward is added.

The logit response function analogous to (3.1.2) is again a ratio of exponential functions of expected payoffs, with the sum in the denominator being replaced by an integral, which is a constant (independent of x) that forces the resulting density to integrate to 1:

$$f(x) = \frac{e^{\lambda\pi(x)}}{\int_{\underline{x}}^{\bar{x}} e^{\lambda\pi(y)} dy} \quad (7.2)$$

As with the discrete case, the expected payoff and quantal response functions jointly determine the equilibrium (McKelvey and Palfrey, 1996). In particular, the equilibrium requirement is that the $f(x)$ density representing beliefs in the expected payoff function (3.7.1) must match the choice density determined by the stochastic response function (3.7.2).

In order to obtain a single differential equation in the equilibrium density, differentiate both sides of (3.7.2) with respect to x and then use (3.7.2) to simplify the result:

$$f'(x) = \lambda \pi'(x) f(x), \quad (7.3)$$

where the derivative of expected payoff function (3.7.1) for the Traveler's Dilemma is:

$$\pi'(x) = -2Rf(x) + [1 - F(x)]. \quad (7.4)$$

Equation (3.7.3), which will be called the *logit differential equation*, is a general condition that holds for the equilibrium in all logit continuous models, given the expected payoff derivative for the model being considered. For the Traveler's Dilemma, we can combine these two equations to obtain a single differential equation in the logit equilibrium density:

$$f'(x) = \lambda \left(-2Rf(x) + [1 - F(x)] \right) f(x), \quad (7.5)$$

For specific parameter values, this equation can be solved numerically, given the boundary condition that the density integrates to 1, but our interest is in determining the general effects of a change in the penalty/reward parameter on the equilibrium. The following proposition is phrased somewhat more generally so that it can be used to evaluate logit QRE predictions for some of the other models to be considered. All of these models will have the property that the expected payoff derivative can be expressed as a function of a payoff parameter, α , and the decision variable x , both directly and indirectly through $F(x)$ and $f(x)$ functions. Thus we can write the expected payoff derivative as $\pi'(\alpha, F, f, x)$, where the final argument pertains only to the direct effect. For example, the derivative in (3.7.4) can be expressed: $\pi'(\alpha, F, f, x) = -2\alpha f + [1 - F]$, which is a special case with no direct x effect.

Proposition 3.4. Consider an n -person game in which each player must select a decision, x , from an interval $[\underline{x}, \bar{x}]$, and let the symmetric logit equilibrium distribution and density functions be denoted by $F(\cdot)$ and $f(\cdot)$. If the expected payoff derivative with respect to a player's decision can be expressed a function of x , $F(\cdot)$, $f(\cdot)$, and a payoff parameter α , and if $\partial\pi'_i(\alpha, F, f, x)/\partial x \leq 0$, then an increase in α will lower the equilibrium distribution in the sense of first-degree stochastic dominance if $\partial\pi'_i(\alpha, F, f, x)/\partial\alpha < 0$, and an increase in α will raise equilibrium distribution if $\partial\pi'_i(\alpha, F, f, x)/\partial\alpha > 0$

Proof (Anderson, Goeree, and Holt, 2002). Let α and α^* be the two parameters, with $\alpha^* > \alpha$, and let the corresponding equilibrium distributions be represented by $f(x)$ and $f^*(x)$, with analogous notation for the distribution functions. First, consider the case where the expected payoff derivative is decreasing in α . We want to show that the higher parameter (α^*) generates lower claims in the sense that $F^*(x) \geq F(x)$ for $x \in [\underline{x}, \bar{x}]$. To obtain a proof by contradiction, suppose instead that $F^*(x) < F(x)$ for some region of claims, as shown in Figure 3.10. Consider the horizontal line in the figure, which is drawn at the level, y , of maximum horizontal difference. The associated claims, x and x^* , yield the same values of their respective cumulative distributions: $F(x) = F^*(x^*) = y$. Since the logit equilibrium distribution functions will be continuous and differentiable, the slopes will be equal at the point of maximum horizontal difference: $f(x) = f^*(x^*)$. The second-order necessary condition for the horizontal distance to be maximized is that $f^{*'}(x^*) \geq f'(x)$, i.e. that the slope of the $F^*(x)$ be increasing more rapidly at x^* . The slope of the density at x is:

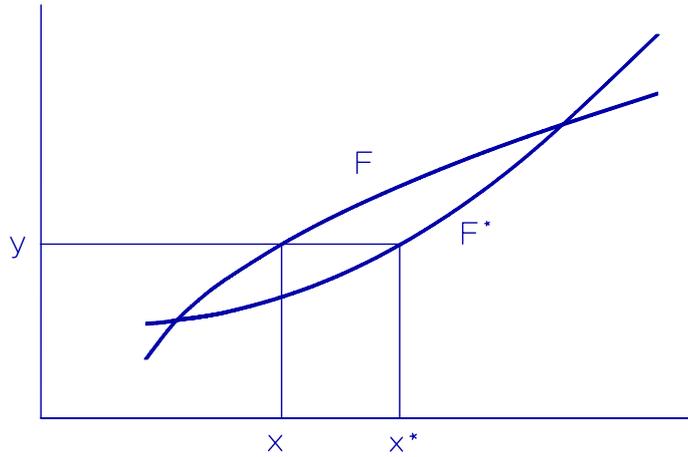


Figure 10: A *Contradiction*.

$$f'(x) = \lambda \pi'_i(\alpha, F(x), f(x), x) f(x),$$

and the slope at x^* is determined by the same function,

$$f^{*'}(x^*) = \lambda \pi'_i(\alpha, F^*(x^*), f^*(x^*), x^*) f(x^*).$$

By construction, $F(x) = F^*(x^*)$ and $f(x) = f^*(x^*)$, but with $\alpha^* > \alpha$ and $x^* > x$, where both inequalities reduce the value of expected payoff derivative by assumption. It follows that $f'(x) > f^{*'}(x^*)$, which violates the second-order necessary condition for the horizontal difference to be maximized, as stated above. The other case, where the expected payoff derivative is increasing in α and $\alpha^* < \alpha$, is proved in an analogous manner. QED

Since the expected-payoff derivative for the Traveler's dilemma game is not a function of x and is decreasing in R , the proposition applies and an increase in the penalty/reward parameter will reduce the claim distribution in the sense of first-degree stochastic dominance.

8. A Model of Imperfect Price Competition

As with the Prisoner's Dilemma, the Traveler's Dilemma provides a paradigm designed to make an important point about strategic interactions in an intentionally simplified setting with minimal institutional detail. The minimum-claim feature of the Traveler's Dilemma,

however, arises naturally in many types of economic models, especially those involving price competition. Indeed, Camerer (2003) discusses the Traveler’s Dilemma in a market context. In this section, we summarize the earlier analysis of price competition in an experimental setting designed to evaluate the “out-of-sample” predictive accuracy of the logit equilibrium model that was estimated for the Traveler’s Dilemma data.

Capra *et al.* (2002) report a laboratory experiment based on duopoly with simultaneous price choices from a range $[\underline{P}, \bar{P}]$. The lower of the two prices will be denoted by P_{min} , and the firm selecting this price will, naturally, obtain a larger market share. Demand is perfectly inelastic at a quantity of $1 + \alpha$, with the sales for the low-price firm normalized to 1, and the sales for the high-price firm being α units, with $\alpha < 1$. In the event of a tie, the demand is shared equally, with each firm selling $(1 + \alpha)/2$ units. The product is homogeneous, so the high-price firm has to meet the lower price to sell any units, but the process of meeting the lower price results in some loss of sales.¹⁹ For simplicity, there are no costs, so the low-price firm earns P_{min} and the other firm, with sales of α , earns αP_{min} .

As long as the market share of the high-price firm is smaller, i.e. $\alpha < 1$, the usual Bertrand logic applies and either seller would have an incentive to undercut any common price. In addition, a unilateral price increase above a common level would be unprofitable because it would reduce market share and would not alter the minimum price. Therefore, the unique Nash equilibrium for this game is the lowest price, \underline{P} .²⁰ To summarize, the Nash price is the lowest price when α is below one. When $\alpha = 1$, however, the market shares are equal, and each firm has a weakly dominant strategy of charging the maximum price.

The starkly competitive nature of the assumed price competition in a Nash equilibrium is implausible if one expects that the degree of buyer inertia to affect pricing behavior. When $\alpha = 0.9$, for example, there is little lost in the way of sales from having the high price, a firm might be willing a risk price increase if there is some small chance that the other firm will do the same. In contrast, lower prices might result from lower values of α . The motivation behind the design was the expectation that average price might be a smoothly increasing in the degree of buyer inertia, parameterized by α , instead of staying at the Nash prediction and then jumping sharply to the upper limit when $\alpha \geq 1$.

The experiment was run at the University of Virginia with six groups of 10 subjects each. Participants were randomly paired for a sequence of 10 rounds, and in each round they would select prices simultaneously from an interval $[\$0.60, \$1.60]$. Instead of using a wide range of buyer inertia parameter values in different sessions, as in Capra *et al.* (1999), the authors ran multiple sessions in each of two treatments: three sessions with $\alpha = 0.2$, and

¹⁹For example, if some of the buyers were covered by “meet-or-release” clauses, then these buyers would not be able to switch without asking the high-price firm to meet the lower price, which it would do to avoid losing all sales. A large fraction of the other buyers (who do not have meet-or-release clauses in purchase contracts) would presumably go straight to the firm with the lower price, and for this reason, the firm with the lower initial price would have higher sales. See Holt and Scheffman (1987) for an analysis of the effects of meet-or-release and other best-price provisions on market outcomes.

²⁰It can be shown that there is no mixed-strategy Nash equilibrium as long the upper bound \bar{P} is finite. If demand were inelastic at any price, no matter how high, then a mixed-strategy Nash equilibrium exists, but it has unintuitive comparative-statics properties, i.e. an increase in α (the sales for the high-price firm) would result in *lower* prices on average.

three sessions with $\alpha = 0.8$.

The average prices in penny amounts for the final 5 rounds of each session are shown in Table 3.6. These averages are higher for each of the three high- α sessions than is the case for each of the three low- α sessions, despite the fact that the Nash equilibrium is \$0.60 for all 6 sessions. A simple non-parametric test can be constructed by noting that there are 20 different ways that 6 numbers from two categories can be ranked, i.e. “six choose three” $= 6!/(3!3!) = 20$. Each of these 20 outcomes is equally likely under the null hypothesis of no effect, but the one shown in the table is the most extreme in the direction of the treatment effect. Therefore, the null can be rejected with a 5 per cent (1/20) level of significance.

Table 3.6 Average prices (standard deviations) and logit predictions for $\lambda = 1.2$.

(Source: Capra *et al.*, 2002)

	Sessions 1, 2	Sessions 3, 4	Sessions 5, 6	Pooled	Logit Predictions
Low α treatment:	63 (14)	72 (20)	73 (32)	69 (13)	78 (7)
High α treatment:	102 (14)	126 (31)	134 (17)	121 (13)	128 (6)

For a given value of λ , equations (3.1.1)-(3.1.3) can be used to obtain a set of 101 simultaneous equations for the logit equilibrium choice probabilities associated with each penny amount on the interval from 60 to 160 cents, and these equilibrium probabilities are then used to calculate the expected values and standard deviations (for a sample of size 10) of the equilibrium price choices. With $\lambda = 1.2$, which was the value estimated from the Traveler’s Dilemma experiment, the logit predictions are shown in the far right column of Table 3.6. Notice that these predictions track the average prices “out-of-sample” for each treatment, shown in the adjacent column just to the left of the predictions. In fact, this theoretical analysis was done *before the experiment was run*, and the particular values of α (0.2 and 0.8) were selected to make it likely that the price averages would converge to levels on opposite sides of the midpoint of the range of feasible prices.²¹

Comparative Statics

The numerical analysis of the effects of the treatment change for this game only applies to specific parameter values, so it is useful to obtain a more general result. The expected payoff function is:

$$\pi(p) = \int_p^P \alpha y f(y) dy + p[1 - F(p)], \quad (8.1)$$

The first term on the right is an integral over prices below the player’s own price, p , which are multiplied by α , the sales quantity when the player’s price is high. The second term is

²¹In general, the value of the precision parameter λ need not be the same for different groups of subjects in different strategic situations, since this parameter will depend on the degree of randomness in individual decisions, which may be influenced by the difficulty of the decision task, confusion about payoffs, the cognitive abilities of the subjects, and any heterogeneity in various non-monetary motivations of the participants. In this case, however, the some of the key aspects of the experiment (subject pool, matching protocol, payoff levels, and complexity of the game) were essentially the same as with the Traveler’s Dilemma experiment discussed above.

the product of the player's price and the probability that the other's price is higher (times the sales quantity, which has been normalized to equal 1). The derivative of this expected payoff is:

$$\pi'(p) = [1 - F(p)] - (1 - \alpha) p f(p). \quad (8.2)$$

The logit differential equation in (3.7.2) becomes:

$$f'(p) = \lambda \left([1 - F(p)] - (1 - \alpha) p f(p) \right) f(p), \quad (8.3)$$

This derivative is a decreasing function of the price choice, which corresponds to the choice variable x in Proposition 3.4, so that proposition also applies in the present case. In particular, the expected payoff derivative is an increasing function of α , so an increase in the market share parameter will increase prices in the sense of first-degree stochastic dominance.

Learning and Convergence to Equilibrium

Keynes once offered a sharp criticism of equilibrium analysis ("in the long run, we are all dead"). In many situations, particularly in macroeconomics, the patterns of price adjustment and directions from which prices converge may be of independent interest. Figure 3.5 shows the average prices, by round, for each of the six sessions, with the high-buyer-inertia sessions shown by the dashed lines at the top. In addition, the thick line shows the average over all sessions in a treatment. Prices for the high-buyer-inertia sessions start high and stay high, with no tendency to approach the Nash prediction of \$0.60. In contrast, the prices in the low-buyer-inertia sessions (with $\alpha = 0.2$) start at an intermediate level and tend to decline slowly towards the Nash prediction. Obviously, these dynamic patterns cannot be explained by an equilibrium model, so we must consider models of learning and adjustment. These models will be developed in detail in a later chapter, but it is useful to provide a preview here, in order to clarify the distinction between equilibrium and learning models.

To obtain a tractable learning model, let the 101 penny amounts on the range [\$0.60, \$1.60] be indexed by j . Consider a specific person (with index i , which will be suppressed), and let this person's beliefs about the other's price be represented by 101 probabilities that sum to 1; these probabilities will be denoted by p_j . In the empirical implementation of this model, the initial beliefs are uniform, so $p_j = 1/101$ for $j = 1, \dots, 101$. An easy way to generalize this setup is to give each of the 101 prices a "weight" of 1, and then let the belief corresponding to a particular price be its weight divided by the sum of the weights for all prices. Then if a particular price is observed, its weight can be increased, which will increase the belief probability associated with that price. One way to do this is to add 1 to the weight of a price that is observed. Specific prices that are observed recently may have a larger impact on beliefs than prices observed in prior rounds, and one way to deal with such "recency effects" is to discount all weights based on past observations by a factor, ρ , with $0 < \rho < 1$. Thus a price with a weight w that is observed in a particular round would have its weight increased to $\rho w + 1$, and all other prices would have their weights decreased from w to ρw . Thus a price observed t rounds in the past would have the weight of that

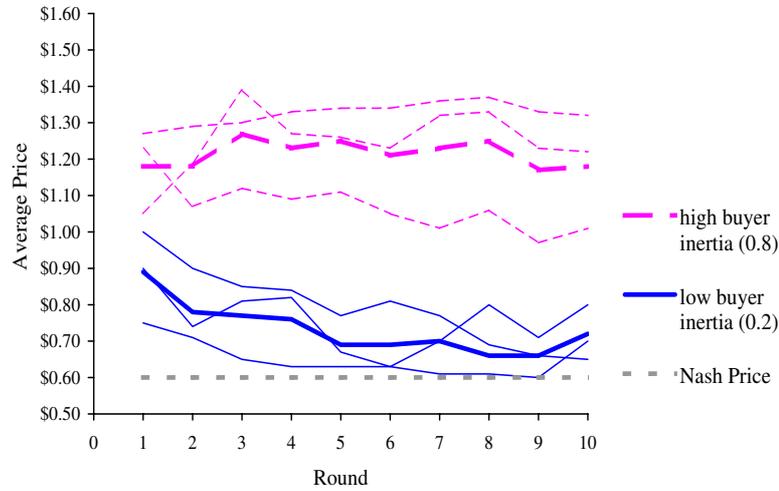


Figure 11: *Average price data.* Key: *thin dashed lines are for 3 sessions with $\alpha = 0.8$, thin continuous lines are for 3 sessions with $\alpha = 0.2$, thick dashed and continuous lines are averages pooled across all 3 sessions for a given treatment.* (Source: Capra, et al. 2002)

observation reduced by a factor of ρ^t , which yields a model of “geometric fictitious play.”²² In the limit, if $\rho = 1$ there is no discounting and each price observation has the same impact on beliefs, regardless of how long ago it was observed.

To summarize, the simplest geometric fictitious-play learning model begins with uniform initial belief probabilities for each price that are updated as the individual observes specific prices selected by others in a sequence of rounds. At a specific point in the experiment, the beliefs are represented by probabilities p_j for the price with index j , and these beliefs about the other seller’s price can be used to calculate the expected payoff for selecting the price with index j , which will be denoted by π_j , as before. These expected payoffs, one for each possible price choice, in turn can be used to determine choice probabilities via the logit response function in equation (3.1.2). In early rounds, the beliefs determined by the learning rule will typically not match the choice probabilities determined by the logit response function (3.1.2), i.e. the equilibrium consistency of choice and belief probabilities in equation (3.1.3) is *not* imposed in a learning model.

Note that these calculations must be made for specific values of the precision parameter, λ , that is used in the logit response function, and for the recency parameter, ρ , used in the learning rule. Then specific parameter values can be used to *simulate* the learning

²²Cheung and Friedman (1997) provide a test of the geometric fictitious play model for 2x2 games, and significant generalizations of this model are discussed in Camerer (2003).

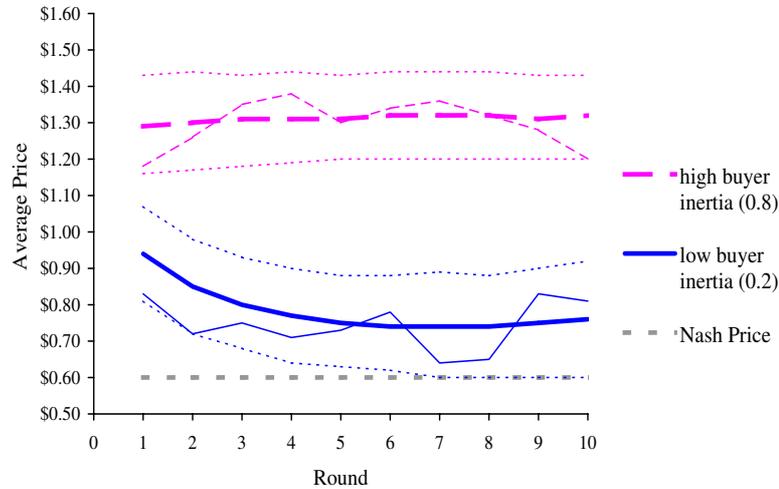


Figure 12: *Average prices for simulated sessions. Key: The thick dashed line is for 1000 simulations with a high α , and the thin dashed line is for a single representative simulation. The thick continuous line is for 1000 simulations with a low α , and the thin continuous line is for a single representative simulation. The dotted lines show the range around the overall averages (plus or minus two standard deviations). (Source: Capra, et al. 2002)*

process, by setting up a group of 10 simulated players, each with with equal initial weights for each price(a uniform prior). The resulting beliefs and expected payoffs determine the choice probabilities, and actual choices can be randomly generated. Each simulated player is randomly matched with another, and the partner’s simulated decision is used to update each players beliefs before the next set of simulated prices are randomly generated.

Capra et al. (2002) used the Traveler’s Dilemma data from an earlier paper to obtain an estimate of 0.75 for ρ , and then ran a set of 1000 simulated sessions for each treatment of the price competition game. A simulated session involved 10 players, who were randomly matched for 10 rounds. With $\alpha = 0.2$, the price averages start at about \$0.93 and then decline before levelling off at about \$0.75, as shown by the thick lower line in Figure 3.12. For the high- α treatment, simulated prices start high, at about \$1.30, and stay at that level for all rounds, see the thick dashed line at in the upper part of the figure. The smooth dotted lines show the range of plus or minus two standard deviations for the simulated sessions for each treatment. The thin kinked lines (dashed for $\alpha = 0.8$ and solid for $\alpha = 0.2$) show the average prices for a typical result of a single simulated session for each treatment.

Notice that these two individual simulated sessions shown in Figure 3.12 have price sequences that do not completely settle down in later rounds, as is also the case for for human subjects shown in Figure 3.11. Continuing randomness is introduced by the logit

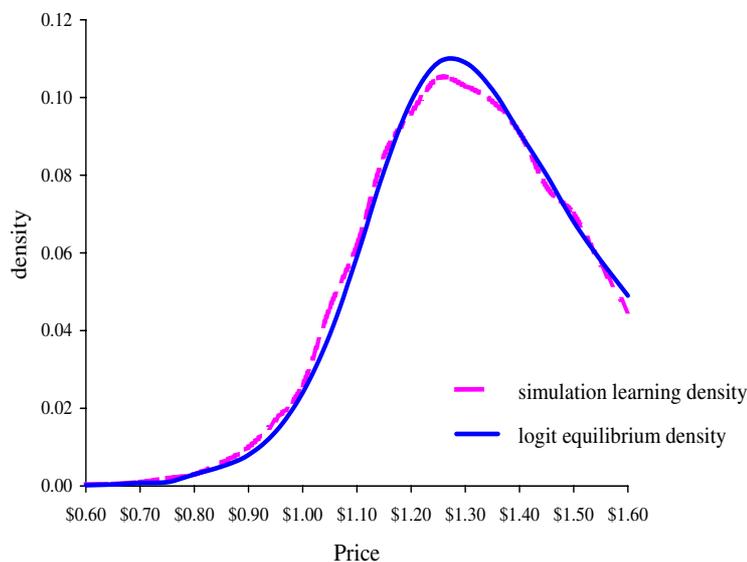


Figure 13: *A comparison of the logit equilibrium density with the density resulting from 1000 simulations of the geometric learning model with a recency parameter of $\rho = 0.75$. (Source: Capra et al. 2002)*

response functions, but there may be a second factor at work in the simulations, since the recency parameter will magnify the effects of more recent observations. This would cause individuals to have different belief distributions, even after very long histories of play. Since all belief distributions are identical in a symmetric quantal response equilibrium, it will be the case that heterogeneity caused by heterogeneous “histories” across simulated subjects will produce a steady state choice distribution that is slightly “flatter” than the logit quantal response equilibrium distribution. As can be seen from Figure 3.13, this difference is quite minor for the imperfect price competition game under consideration. For simulations of the imperfect price competition game with $\rho = 0.75$, the simulated learning and logit equilibrium distributions are virtually identical. Stronger recency effects, however, would produce greater differences. These ideas are developed in more detail in a later chapter on learning, where we introduce the idea of a quantal response learning equilibrium, or a steady state distribution based on noisy responses to current beliefs.

9. Minimum Effort Coordination Games

Coordination problems arise naturally in many economic activities that involve joint production or other interactions with positive externalities. For example, if the assembly

of a product involves two components that are made by different individuals, then excess production by one person can be costly in the sense that it does not increase the number of assembled units completed. In this case, the output of the joint production process is determined by the minimum of the two production efforts, which results in a need to coordinate. Economists have long been interested in coordination games, given the possibility that coordination failure may result in unfavorable outcomes for all. After the Prisoner's Dilemma, the coordination game is perhaps the most widely studied paradigm in game theory.

In theoretical models with multiple, Pareto-ranked Nash equilibria, it used to be common to *assume* that players could somehow coordinate on the equilibrium that is best for all. This assumption might be reasonable when there is well-institutionalized coordination device or norm, but laboratory experiments have shown that such coordination is not to be expected uniformly, especially with large numbers of players. In particular, Van Huyck *et al.* (1990) found convergence to the Nash equilibrium that is *worst* for all, at least in sessions with large numbers of players (about 15).

In chapter 2, we introduced a simple 2x2 minimum-effort game with just two possible effort levels, 1 and 2, and with Nash equilibria at both low and high common effort levels (see Table 2.5). Van Huyck *et al.* ran their experiments with a somewhat larger (7x7) coordination game, with seven Pareto-ranked Nash equilibria, one at each common effort level. In this section we will consider a minimum-effort game with a continuum decisions, i.e. in which each person selects an effort, e_i , from the interval $[e, \bar{e}]$. Payoffs determined by the minimum effort minus the cost of the person's own effort:

$$\pi_i(e_1, \dots, e_n) = \min\{e_1, \dots, e_n\} - ce_i, \quad i = 1, \dots, n, \quad (9.1)$$

where $0 < c < 1$ is an effort cost parameter that is the same for all players. This cost assumption ensures that any common effort is a Nash equilibrium for this game, since a unilateral decrease will reduce the minimum by 1 but will only result in a cost savings of $c < 1$. Conversely, a unilateral effort increase above a common level will raise a player's cost but will not raise the minimum. Since all of these equilibria are strict, standard refinements cannot be used to select one of them. More worrisome is the fact that the set of equilibria is not altered by non-critical changes in the the cost of effort or the number of players, which contradicts the intuition that coordination on high-effort outcomes might be more difficult with more participants and a higher cost of effort.²³

The Nash equilibrium analysis in the previous paragraph began by considering deviations from a known common effort level. This raises the question of how such a level would be known if there are multiple equilibria. The payoff parameters, which affect the costs of deviation in each direction, would presumably have some impact on what players think others will do. Note that an increase in effort reduces one's payoff by c , and a downward

²³Anderson, Holt, and Goeree (2001) showed that there is also a continuum of mixed-strategy Nash equilibria in which players randomize between two specific effort levels. These equilibria, however, have the perverse property that an increase in the cost of effort results in an *increase* in the probability of the high-effort decision.

deviation reduces the minimum by 1 and the cost by c , so the cost of a downward deviation from a known common effort is $1 - c$. The costs of deviating are the same in each direction if $c = 1 - c$, i.e. $c = 1/2$, and the payoff loss from an increase in effort is greater than the payoff loss from a decrease if $c > 1/2$. These quick calculations at least provide some intuition for why effort levels might be negatively correlated with costs.

Recall that the “*risk-dominant*” equilibrium for the 2x2 coordination game in Table 2.5 (in the previous chapter) involved low efforts if the effort cost was less than $1/2$ and high efforts otherwise. There is no widely accepted way of generalizing the notion of risk dominance to games with more than two feasible decisions, but there is a related and somewhat more general notion, the *maximization of potential*, which does correspond to risk dominance in the 2x2 case. The general idea here is to find a function, which when maximized, yields a Nash equilibrium. The hope is that the maximization of potential will help select among multiple Nash equilibria in games where other selection devices fail. Consider an n -person normal-form game, where player i selects a decision x_i from some interval. Let the players’ payoff functions be denoted $\pi_i(x_1, \dots, x_n)$, $i = 1, \dots, n$. This game is a *potential game* if there exists a potential function of all players’ decisions, $V(x_1, \dots, x_n)$, with the property that the partial derivatives of the potential function match the partial derivatives of the payoff functions with respect to the players’ own decisions:

$$\frac{\partial V(x_1, \dots, x_n)}{\partial x_i} = \frac{\partial \pi_i(x_1, \dots, x_n)}{\partial x_i} \text{ for } i = 1, \dots, n. \quad (9.2)$$

where the payoff derivatives exist.²⁴ The class of potential games includes a well known variety of interesting games.²⁵

The potential function for the minimum-effort game in (3.9.1) is:

$$V(e_1, \dots, e_n) = \min\{e_1, \dots, e_n\} - \sum_{i=1}^n c e_i, \quad (9.3)$$

where the inclusion of the effort costs for all players on the right side is needed to ensure that the partial derivatives of the potential function match the corresponding payoff function derivatives. The maximization of this potential function will require that all efforts be equal (to avoid wasted excess effort), so the potential can be expressed as a function of a common effort, e : $V = e - nce$. The derivative of this expression is $1 - nc$, so it is maximized at the lowest possible effort \underline{e} when $c > 1/n$ and at the highest effort \bar{e} when $c < 1/n$. Thus the critical value of c is $1/2$ when there are two players, and an increase in the number of players makes it harder to coordinate on the highest effort outcome in the sense that the relevant range of costs is reduced.²⁶ Nevertheless, the prediction based on the maximization

²⁴Rosenthal (1973) introduced the notion of a potential game, although some others had previously used the general idea in the analysis of oligopoly problems.

²⁵This class includes all 2x2 games, and some versions of public goods and Cournot oligopoly games. See Monderer and Shapley (1996) and Anderson, Goeree, and Holt (2004) for more examples.

²⁶It can be shown that the maximization of potential for the 2x2 coordination game in Table 2.5 is at low efforts when $c > 1/2$ and at high efforts when $c < 1/2$, so in this case, the predictions based on potential

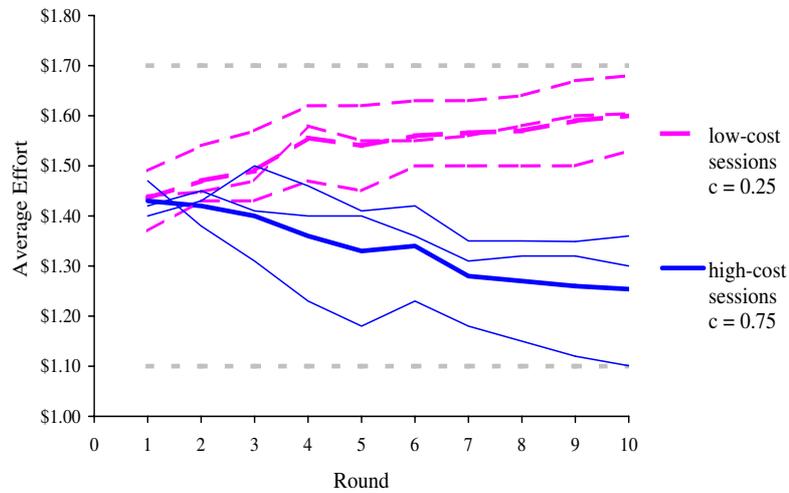


Figure 14: *Average effort data for a coordination game experiment. Key: thin dashed lines are for 3 sessions with $c = 0.25$, thin continuous lines are for 3 sessions with $c = 0.75$, and thick lines, dashed and continuous, are averages pooled across all 3 sessions for a given treatment. (Source: Goeree and Holt, 2004)*

of potential is at one of the extreme ends of the set of feasible efforts, unless c is exactly $1/n$. Intuitively, one might expect more of a smooth transition from low to high effort levels as the cost of effort is increased in the neighborhood of the critical point. The QRE approach accomplishes this by the introduction of noisy equilibrium behavior.²⁷

The invariance of Nash predictions to changes in the cost of effort was the motivation for several coordination game experiments reported in Goeree and Holt (2004). They began with a simple 2-person minimum-effort coordination game, with an effort range from \$1.10 to \$1.70. Subjects were recruited in groups of 10, and they played a series of 10 rounds with random matching. Since the critical point (based on maximum potential) was $1/2$ for these two-person games, the experimental design involved cost values, $1/4$ and $3/4$, that were not too close to $1/2$ and not too extreme either. Three sessions were run under either treatment condition.

The round-by-round average effort decisions for these two treatments are plotted in Figure 3.14, where the thin lines represent individual sessions and the thick lines represent the averages pooled for all sessions in a treatment. Notice that the average efforts for both and risk-dominance coincide (Goeree and Holt, 2004).

²⁷A more detailed discussion of the relationship between logit equilibria and potential maximization will be provided in Chapter 6, where we introduce the notion of a *stochastic potential function*, which has extreme points that are logit equilibria.

treatments begin at about the midpoint of the range of feasible effort decisions, and the null hypothesis that the effort distributions for the two treatments are equal cannot be rejected at the 5 percent level using a standard Kolmogorov-Smirnov test. The average effort series, however, show a clear separation by the fifth round, and for the final five rounds, the null hypothesis of no treatment effect can be rejected at the 5 percent level.²⁸

The QRE analysis of this game will be based on the logit differential equation (3.7.3), so we need a formula for the expected payoff derivative as a function of the distribution of the other player's effort choices. Let the distribution function for a player's effort decisions be $F(e)$, with density $f(e)$. One way to proceed would be to use this notation to obtain an expression for the expected value of the payoff function in (3.9.1), and then take the derivative with respect to a player's own effort, e_i . A quicker and more intuitive way to get the same result is to note that the cost of effort is linear in effort, so there will be a term in the expected payoff derivative that is the negative of the cost parameter. In addition, an increase in effort e_i increase will raise the minimum if all $n - 1$ other players' efforts are above e_i , which occurs with probability $[1 - F(e_i)]^{n-1}$ so the derivative of the expected payoff is:

$$\pi'(e) = -c + [1 - F(e)]^{n-1}, \quad (9.4)$$

This derivative is substituted into the formula for the logit differential equation (3.7.3) to obtain:

$$f'(x) = \lambda \left(-c + [1 - F(e)]^{n-1} \right) f(x). \quad (9.5)$$

Note that the expected payoff derivative is decreasing in both c and n . It follows from Proposition 3.4 that an increase in either the effort cost or the number of players will decrease the distribution of efforts in the sense of first-degree stochastic dominance.

The theoretical comparative-statics analysis indicates that the *qualitative* predictions of the logit analysis are consistent with the data patterns, but the quantitative dimensions of the treatment effect are explained as well. Table 3.7 contains the average effort levels for each session, for rounds 6-10, with standard deviations in parentheses. The averages for the pooled data are 159 for the low-cost sessions and 126 for the high-cost sessions. The logit predictions were obtained by dividing the set of feasible effort choices into one-penny intervals and using the logit response probabilities in equation (3.1.2) to estimate a precision parameter of $\lambda = 1.35$. As before, the resulting equilibrium effort distribution can be calculated and used to compute the expected value of the equilibrium effort for each treatment level. These predictions, shown in the far-right column of Table 3.7, are quite close to the corresponding averages for the pooled data.

²⁸The intuition behind the test, as explained previously, is that there are "six choose three" = 20 possible ways that the average efforts for the six sessions could have been ranked by treatment, and of these, the most extreme ranking was observed, i.e. with all three low-cost sessions producing the top three effort averages. Hence the probability of this outcome under the null hypothesis is 1/20, or 5 percent.

Table 3.7 Average efforts (standard deviations) and logit QRE predictions for $\lambda = 1.35$.
(Source: Goeree and Holt, 2004)

Cost	1st Session	2nd Session	3rd Session	Pooled	QRE Predictions
Low ($c = 1/4$)	151 (10)	166 (5)	159 (12)	159 (11)	154 (12)
High ($c = 3/4$)	131 (11)	112 (5)	135 (11)	126 (14)	126 (12)

Finally, consider the interesting patterns of adjustment in Figure 3.14, with a more or less symmetric spreading out of decisions from the midpoint. Goeree and Holt (1999) report estimates of a geometric learning model for these data, and simulations that reproduce this pattern.

10. The All-Pay Auction

Many allocation mechanisms have the property that a prize is awarded on the basis of costly activities of potential recipients. For example, exclusive licenses might be awarded to the person who mounts the most intensive lobbying effort (Tullock, 1967), or to those who wait in line the longest (Holt and Sherman, 1982). Tullock (1980) initiated a large literature in which it is assumed that the probability of obtaining the prize is an increasing function of the lobbying effort. The limiting case, in which the award is always made to the person with the maximum effort, is known as an “all-pay auction,” since the losing contenders are not reimbursed for their efforts. An auction-like selection mechanism is appropriate for cases in which the efforts have some value to the person making the allocation or cases in which the efforts are observed, and rewarding high effort is considered to be fair, e.g. standing in line. Since lobbying activities entail social costs, the main focus of this literature is on the aggregate costs, which tend to “dissipate” the net social value of the prize being awarded.

The efforts will be referred to as bids, so the prize goes to the highest bidder, but all must pay their bid amounts. The simplest case is one in which the prize has with a known value of V that is the same for all. Each of the n bidders make simultaneous bids, b_i for $i = 1, \dots, n$, which must be in an interval $[0, B]$, where it is assumed that $B \geq V$. The cost of a bid of b_i is cb_i , where $c > 0$. In many contexts, the parameter c might be set to 1.

We begin with a derivation of the Nash equilibrium, which will serve as a benchmark for QRE comparisons. If others’ bids are known and are less than V , then the best response would be to bid slightly above the others. Hence, the Nash equilibrium will involve randomization. In a symmetric Nash equilibrium with a common bid distribution denoted by $F^*(b)$, the probability of winning with a bid of b is $F^*(b)^{n-1}$, and the expected payoff is:

$$VF^*(b)^{n-1} - cb. \quad (10.1)$$

Since exit, with a zero payoff, is always an option, the expected payoff in a mixed-strategy Nash equilibrium must be non-negative. This implies that the lowest bid over which randomization occurs must be 0, which generates an expected payoff of zero. All bids over

that are selected with positive probability will, therefore, have zero expected payoffs, and the Nash distribution can be found by equating the expected payoff in (3.10.1) to 0, to obtain: $F^*(b) = (cb/V)^{1/(n-1)}$ for $b \in [0, V/c]$, and a density of $f^*(b) = \frac{1}{n-1} \left(\frac{c}{V}\right)^{1/(n-1)} b^{(2-n)/(n-1)}$. Since the distribution function is increasing in both c and n and decreasing in V , it follows that bids will be increased in the sense of first-degree stochastic dominance when there is a reduction in cost c or the number of bidders, or when there is an increase in the prize value, all of which are intuitive properties. Notice that the Nash equilibrium density is constant (uniform) for $n = 2$ and decreasing in b for $n > 2$.²⁹

Next consider the symmetric logit QRE, which will be characterized by the logit differential equation (3.7.3). First, note that the expected payoff derivative is:

$$\pi'(b) = V(n-1)F(b)^{n-2}f(b) - c, \quad (10.2)$$

which is substituted into the logit differential equation to obtain:

$$f'(b) = \lambda [V(n-1)F(b)^{n-2}f(b) - c] f(b), \quad (10.3)$$

The expected payoff derivative (3.10.2), viewed as a function of F , f , and b does not depend on b directly, only indirectly through the distribution and density functions, so Proposition 3.4 applies. It follows that an increase in V or a decrease in c will raise the QRE bid distribution in the sense of first-degree stochastic dominance. The numbers effect cannot be determined by this method, but Anderson, Goeree, and Holt (1998a) found an explicit solution to (3.10.2), which has the property that an increase in n will lower the equilibrium distribution of bids.

A Symmetric Example with Two Bidders

If $n = 2$, the logit differential equation in (3.10.3) reduces to $f' = \lambda (Vf - c) f(b)$, which can be solved explicitly. First, divide both sides by $(Vf - c)f$ and rewrite the result:

$$\frac{f'}{(Vf - c)f} = \frac{1}{c} \left(\frac{Vf'}{Vf - c} - \frac{f'}{f} \right) = \lambda, \quad (10.4)$$

This equation can be integrated to obtain:

$$f(b) = \frac{c}{V[1 - K \exp(cb\lambda)]}, \quad (10.5)$$

where k is a constant of integration obtained by equating the integral of the density over $[0, B]$ to one:

$$K = \frac{\exp(\lambda[V - cB]) - 1}{\exp(\lambda V) - 1}. \quad (10.6)$$

²⁹Baye *et al.* (1996) show that this is the unique Nash equilibrium for this formulation of the all-pay auction.

It can be verified that this density is positive on $[0, B]$. If the upper bound is exactly equal to V/c , so that bidding for a sure loss is ruled out, then $K = 0$ and the density in (3.10.5) is constant on its support, independent of λ . The resulting uniform distribution is the same as the Nash equilibrium distribution (with $n = 2$) computed previously. If overbidding errors are allowed, i.e. $B > V/c$, then the constant, K , determined by (3.10.6) is negative, and (3.10.5) implies that the logit equilibrium density is everywhere decreasing on its support.³⁰

An Asymmetric Example with Two Bidders: Own-Payoff Effects

Next we consider the case of value asymmetries, which can lead to unintuitive comparative statics results for a Nash equilibrium. In particular, an increase in one bidder's value will not alter the Nash density for that player. Suppose that $V_1 > V_2$. With asymmetries, the probability that one person has the high bid is determined by the other person's equilibrium distribution function. Therefore, the expected payoff for bidder 2 is $V_2 F_1(b) - cb$. In a mixed-strategy Nash equilibrium, this expected payoff is constant over the range of randomization, and therefore, F_1 is a function of V_2 and c , but not of V_1 , and vice versa. In contrast, the logit equilibrium does exhibit "own-payoff" effects, as will be demonstrated by solving for the equilibrium densities in the asymmetric, two-bidder model.

In a logit equilibrium, the densities are proportional to exponential functions of expected payoffs: $f_i(x) = K_i \exp^{\lambda(V_i F_j(b) - cb)}$, $i, j = 1, 2, j \neq i$. where K_i is a constant set to ensure that the density integrates to 1. Since a zero bid ensures a payoff of 0 in this game, it follows that $K_i = f_i(0)$, so the equilibrium conditions can be expressed:

$$f_i(x) = f_i(0) \exp^{\lambda[V_i F_j(b) - cb]}, \quad i, j = 1, 2, j \neq i. \quad (10.7)$$

Differentiating the above equations with respect to b produces the logit differential equations:

$$f'_1 = \lambda f_1(V_1 f_2 - c), \quad \text{and} \quad f'_2 = \lambda f_2(V_2 f_1 - c). \quad (10.8)$$

By multiplying the left equation by V_2 and the right equation by V_1 and subtracting, we obtain:

$$V_2 f'_1 - V_1 f'_2 = -\lambda c(V_2 f_1 - V_1 f_2). \quad (10.9)$$

Equation (3.10.9) can be integrated to obtain:

$$V_1 f_2(b) - V_2 f_1(b) = A \exp(-\lambda cb), \quad (10.10)$$

³⁰McKelvey and Palfrey (1995) show that the quantal response equilibrium converges to a Nash equilibrium as the precision parameter, λ , tends to infinity. In the present example with a continuum of strategies, this approach can be used to provide an alternative derivation of the Nash equilibrium bid density, which is uniform on $[0, V/c]$. See Anderson, Goeree, and Holt (1998a, note 12) for details.

where A is a constant of integration. Equation (3.10.10) shows $V_1 f_2$ as a function of $V_2 f_1$ (and vice versa), and these can be substituted into the appropriate equations in (3.10.8) to obtain single differential equation for each of the logit densities, which can be solved explicitly (Anderson, Goeree, and Holt, 1998b, note 15). Our main interest here, however, is on the effects of a change in a player's value on that players' own bid distribution.

Proposition 3.5. *In a two-player all-pay auction, an increase in a player's value results in an increase in that player's logit equilibrium bids (in the sense of first-degree stochastic dominance).*

Proof (Anderson, Goeree, and Holt, 1998b). Suppose that player 1's value increases from V_1 to V_1^* . Equation (3.10.10) can be integrated from 0 to b to obtain: $V_1 F_2(b) = V_2 F_1(b) + (A/\lambda c)[1 - \exp(-\lambda cb)]$. Evaluating this equation at $b = B$ and using the fact that the cumulative distributions are 1, we obtain the constant: $A = (\lambda c)(V_2 - V_1)[1 - \exp(-\lambda cB)]^{-1}$. Using these results, the expected payoff for player 1, $V_1 F_2 - cb$ can be expressed in terms of F_1 , which allows us to rewrite the equilibrium condition (3.10.7) for player 1:

$$f_1(b) = f_1(0) \exp\left(\lambda\left[V_2 F_1 + (V_1 - V_2) \frac{1 - \exp(-\lambda cb)}{1 - \exp(-\lambda cB)} - cb\right]\right). \quad (10.11)$$

If player 1's value is increased from V_1 to V_1^* , the new equilibrium condition is:

$$f_1^*(b) = f_1^*(0) \exp\left(\lambda\left[V_2 F_1^* + (V_1^* - V_2) \frac{1 - \exp(-\lambda cb)}{1 - \exp(-\lambda cB)} - cb\right]\right). \quad (10.12)$$

The structure of the proof will be to show that the distribution functions, F_1 and F_1^* can only cross twice, and hence these crossings must occur at the bids of 0 and B . The final step is to show that F_1 starts out above F_1^* , so that bids are stochastically higher for F_1^* . First, note that at any crossing, $F_1 = F_1^*$, and it follows from (3.10.11) and (3.10.12) that the ratio of slopes at any crossing is:

$$\frac{f_1^*(b)}{f_1(b)} = \frac{f_1^*(0)}{f_1(0)} \exp\left(\lambda(V_1^* - V_1) \frac{1 - \exp(-\lambda cb)}{1 - \exp(-\lambda cB)}\right). \quad (10.13)$$

The right side of (3.10.13) is strictly increasing in b since $V_1^* > V_1$ by assumption, and therefore, the ratio of the densities is increasing at successive crossings. There must only be two crossings, since if there were more, the ratio of densities would either decrease and then increase or the reverse, a contradiction. Since there are only two crossings and the ratio in (3.10.13) is increasing, it must be less than 1 at $b = 0$ and greater than 1 at $b = B$. It follows that $F_1 > F_1^*$ for all interior values of b . Q.E.D.

Rent Dissipation

In order to characterize the extent to which the rent associated with the prize value is dissipated, recall that *ex ante* expected payoffs are zero for all bidders in a Nash equilibrium,

so the rent is fully dissipated. In laboratory experiments, Davis and Reiley (1994) report negative profits and rents that are more than fully dissipated, so the obvious question is whether such excess dissipation can occur in a quantal response equilibrium. It turns out that over-dissipation depends on the number of players, so we return to the case of an n -person all pay auction with symmetric values.

Proposition 3.6. *In a logit equilibrium, there is over-dissipation of rent in the symmetric-value all-pay auction with more than two players and $B \geq V/c$.*

Proof (Anderson, Goeree, and Holt, 1998b). The logit equilibrium density for the symmetric n -person all-pay auction can be written as $f(b) = f(0) \exp(\lambda[V F(b)^{n-1} - cb])$. Recall that the Nash equilibrium density for the symmetric game, $f^*(b)$ is a decreasing function of b when $n > 2$. The ratio of the Nash and logit densities, is:

$$\frac{f(b)}{f^*(b)} = \frac{f(0) \exp(\lambda[V F(b)^{n-1} - cb])}{f^*(b)}. \quad (10.14)$$

The construction of the mixed-strategy Nash equilibrium requires that $V F(b)^{n-1} - cb = 0$, so the term in square brackets on the right side of (3.10.14) is zero when the logit distribution crosses the Nash distribution. Thus $\frac{f(b)}{f^*(b)} = \frac{f(0)}{f^*(b)}$ at the crossings. Since $f^*(b)$ is decreasing, the ratio of densities on the left side of (3.10.14) is increasing at successive crossings. It follows that the ratio is less than 1 at the lower round and greater than 1 at the upper bound, and hence, $F^*(b) > F(b)$ at all interior points. Since the logit distribution lies below the Nash distribution that fully dissipates the rent, the logit distribution produces higher expected bids, which in turn implies that rent is over-dissipated. Q.E.D.

Davis and Reiley (1994) report the results of an experiment in which one of the treatments was an all-pay auction with 4 bidders, $c = 1$, and no upper bound on bids. The implication of Proposition 3.6 is that rents will be over-dissipated. The social costs of rent seeking in the experiment consistently exceeded the prize value, so subjects lost money on average. The losses were more prevalent in early periods. This propensity for losses was dealt with by providing subjects with a relatively large initial cash balance.

11. Conclusion

The applications covered in this chapter span the range from simple matrix and participation games with two decisions to those with a continuum of feasible choices (social dilemma, pricing, and coordination and lobbying games). Despite these differences, the applications covered are similar in certain respects: all have Nash predictions that are unintuitive or

are not affected by changes in key parameters that affect payoffs in asymmetric ways. For example, a high penalty-reward parameter makes it risky to raise claims in the Traveler's Dilemma, just as high effort costs and low degrees of buyer inertia make it more risky to raise decisions (efforts and prices) in the market pricing and coordination games respectively.

Laboratory experiments confirm that the most salient features of the aggregate data are strong behavioral responses to changes in treatment variables that have no effect on the Nash predictions. In each case, these sharp treatment effects are captured by the the relevant discrete or continuous versions of the quantal response equilibrium. These comparative-statics effects are reproduced by using maximum likelihood estimates of the logit precision parameter to compute equilibrium distributions of decisions, and general comparative statics results (for all parameter values) for these models are summarized in Propositions 3.1 - 3.6.

Finally, the interesting dynamic patterns observed in the experiments are evaluated in several ways, both in terms of iterative adjustments of beliefs at the aggregate level, and in terms of individual-specific (geometric) learning models that are estimated and used to simulate round-by-round behavior. These models of learning and adjustment are admittedly crude approximations; more general models will be formulated in Chapter 6.

References

- Anderson, Simon P., Jacob K. Goeree, and Charles A. Holt (1998a) "Rent Seeking with Bounded Rationality: An Analysis of the All-Pay Auction," *Journal of Political Economy*, 106(4), 828-853.
- Anderson, Simon P., Jacob K. Goeree, and Charles A. Holt (1998b) "A Theoretical Analysis of Altruism and Decision Error in Public Goods Games," *Journal of Public Economics*, 70, 297-323.
- Anderson, Simon P., Jacob K. Goeree, and Charles A. Holt (2001) "Minimum-Effort Coordination Games: Stochastic Potential and Logit Equilibrium," *Games and Economic Behavior*, 34(2), 177-199.
- Anderson, Simon P., Jacob K. Goeree, and Charles A. Holt (2002) "The Logit Equilibrium: A Perspective on Intuitive Behavioral Anomalies," *Southern Economic Journal*, 69(1), 21-47.
- Anderson, Simon P., Jacob K. Goeree, and Charles A. Holt (2004) "Noisy Directional Learning and the Logit Equilibrium," *Scandinavian Journal of Economics*, 106(September), 581-602.
- Austin-Smith, D. and J.S. Banks (1996) "Information Aggregation, Rationality, and the Condorcet Jury Theorem," *American Political Science Review*, 90: 34-45.
- Basu, K.(1994) "The Traveler's Dilemma: Paradoxes of Rationality in Game Theory," *American Economic Review*, 84(2): 391-395.
- Baye, Michael R., Dan Kovenock, and Casper G. de Vries (1996) "The All-Pay Auction with Complete Information," *Economic Theory*, 8(2): 291-305.
- Bull, Clive, Andrew Schotter, and Keith Weigelt (1987) "Tournaments and Piece Rates: An Experimental Study," *Journal of Political Economy*,
- Camerer, Colin F. (2003) "Behavioral Game Theory," Princeton, N.J.: Princeton University Press.
- Camerer, Colin F. and D. Lovo (1999) "Overconfidence and Excess Entry: An Experimental Approach," *American Economic Review*, 89(1): 306-318.
- Capra, C. Monica., Jacob K. Goeree, Rosario Gomez, and Charles A. Holt (1999) "Anomalous Behavior in a Traveler's Dilemma," *American Economic Review*, 89(3): 678-690.
- Capra, C. Monica., Jacob K. Goeree, Rosario Gomez, and Charles A. Holt (2002) "Learning and Noisy Equilibrium Behavior in an Experimental Study of Imperfect Price Competition," *International Economic Review*, 43(3): 613-636.
- Cason, Timothy N. and Vai-Lam Mui (2003) "Uncertainty and Resistance to Reform in Laboratory Participation Games," Unpublished draft, Purdue University.
- Davis, Douglas D. and Robert J. Reilly (1982) "Do Too Many Cooks Spoil the Stew? An

- Experimental Analysis of Rent-Seeking and a Strategic Buyer,” *Public Choice*. 95(1-2): 89-115.
- Diekmann, Andreas (1986) “Volunteer’s Dilemma: A Social Trap Without a Dominant Strategy and Some Empirical Results,” in A. Diekmann and P. Mitter, eds., *Procedural Effects of Social Behavior: Essays in Honor of Anatol Rapoport*, Heidelberg: Physica-Verlag.
- Feddersen, Timothy, and Wolfgang Pesendorfer (1996) “The Swing Voter’s Curse ,” *American Economic Review*, 86(3): 408-424.
- Feddersen, Timothy, and Wolfgang Pesendorfer (1998) “Convicting the Innocent: the Inferiority of Unanimous Jury Verdicts Under Strategic Voting ,” *American Political Science Review*, 92: 12-35.
- Feddersen, Timothy, and Wolfgang Pesendorfer (1999) “Elections, Information Aggregation, and Strategic Voting,” *Proceedings of the National Academy of Sciences, U.S.A.*, 96: 10572-10574.
- Frazen, A. (1999) “Group Size and One Shot Collective Action,” *Rationality and Society*, 7: 183-200.
- Goeree, Jacob K. and Charles A. Holt (1999) “Stochastic Game Theory: For Playing Games, Not Just for Doing Theory,” *The Proceedings of the National Academy of Sciences of the United States of America*, 96, 10564-10567.
- Goeree, Jacob K. and Charles A. Holt (2000) “Asymmetric Inequality Aversion and Noisy Behavior in Alternating-Offer Bargaining Games,” *European Economic Review*, 44, 1057-1068.
- Goeree, Jacob K. and Charles A. Holt (2001) “Ten Little Treasures of Game Theory and Ten Intuitive Contradictions,” *American Economic Review*, 91(5), 1402-1422.
- Goeree, Jacob K. and Charles A. Holt (2003a) “Learning in Economics Experiments,” in L. Nadel, ed., *Encyclopedia of Cognitive Science, Vol. 2*, London: Nature Publishing Group, McMillan, pp. 1060-1069.
- Goeree, Jacob K. and Charles A. Holt (2003b) “Coordination Games” in L. Nadel, ed., *Encyclopedia of Cognitive Science, Vol. 2* London: Nature Publishing Group, McMillan, pp. 204-208.
- Goeree, Jacob K. and Charles A. Holt (2004a) “A Model of Noisy Introspection,” *Games and Economic Behavior*, 46(2), pp. 365-382.
- Goeree, Jacob K. and Charles A. Holt (2004b) “An Experimental Study of Costly Coordination,” *Games and Economic Behavior*, forthcoming.
- Goeree, Jacob K. and Charles A. Holt (2005) “An Explanation of Anomalous Behavior in Models of Political Participation,” *American Political Science Review*, forthcoming.
- Goeree, Jacob K. Simon P. Anderson, and Charles A. Holt (1998) “The War of Attrition with Noisy Players” in M. Baye, ed., *Advances in Applied Microeconomics*, 7, Greenwich,

- Ct.: JAI Press, pp. 15-29.
- Goeree, Jacob K., Charles A. Holt, and Susan K. Laury (2002) "Altruism and Noisy Behavior in One-Shot Public Goods Experiments," *Journal of Public Economics*, 83, 257-278.
- Goeree, Jacob K., Charles A. Holt, and Thomas R. Palfrey (2002) "Quantal Response Equilibrium and Overbidding in First-Price Auctions," *Journal of Economic Theory*, 104, 247-272.
- Goeree, Jacob K., Charles A. Holt, and Thomas R. Palfrey (2003) "Risk Averse Behavior in Generalized Matching Pennies Games," *Games and Economic Behavior*, 45, 97-113.
- Guarnaschelli, Serena, Richard D. McKelvey, and Thomas R. Palfrey (2000) "An Experimental Study of Jury Decision Rules," *American Political Science Review*, 94: 407-423.
- Holt, Charles A. (2005) *Markets, Games, and Strategic Behavior*, forthcoming, Boston: Addison-Wesley (**??**): <http://www.people.virginia.edu/~cah2k/expbooknsf.pdf>
- Holt, Charles A., and David Scheffman (1987) "Facilitating Practices: The Effects of Advance Notice and Best-Price Policies," *Rand Journal of Economics*, 18, 187-197.
- Holt, Charles A., and Roger Sherman (1982) "Waiting-Line Auctions," *Journal of Political Economy*, 88, 433-445.
- Kahneman, D. (1988) "Experimental Economics: A Psychological Perspective," in R. Tiez, W. Albers, and R. Selten, *Bounded Rational Behavior in Experimental Games and Markets*, New York: Springer-Verlag, 11-18.
- Luce, Duncan (1959) *Individual Choice Behavior*, New York: Wiley.
- McKelvey, Richard D. and Thomas R. Palfrey (1992) "An Experimental Study of the Centipede Game," *Econometrica*, 60(4), 803-836.
- McKelvey, Richard D. and Thomas R. Palfrey (1995) "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 10(1), 6-38.
- McKelvey, Richard D. and Thomas R. Palfrey (1996) "A Statistical Theory of Equilibrium in Games," *Japanese Economic Review*, 47(2): pp. 186-209.
- McKelvey, Richard D., Thomas R. Palfrey, and Roberto A. Weber (2000) "The Effects of Payoff Magnitude and Heterogeneity on Behavior in 2x2 Games with Unique Mixed Strategy Equilibria" *Journal of Economic Behavior and Organization*, 42: pp. 523-548.
- Meyer, Donald J., John B. Van Huyck, Raymond C. Battalio, and Thomas R. Saving (1992) "History's Role in Coordinating Decentralized Allocation Decisions: Laboratory Evidence on Repeated Binary Allocation Games," *Journal of Political Economy*, 100(2): pp. 292-316.
- Monderer, D. and L.S. Shapley (1996) "Potential Games," *Games and Economic Behavior*, 14: pp. 124-143.
- Ochs, Jack (1990) "The Coordination Problem in Decentralized Markets: An Experiment"

- Quarterly Journal of Economics*, 105(2): pp. 545-559.
- Ochs, Jack (1995) "Games with Unique Mixed Strategy Equilibria: An Experimental Study" *Games and Economic Behavior*, 10: pp. 202-217.
- Palfrey, Thomas R. and Howard Rosenthal (1985) "Voter Participation and Strategic Uncertainty," *American Political Science Review*, 79: pp. 62-78.
- Rosenthal, Robert (1973) "A Class of Games Possessing Pure-Strategy Nash Equilibria," *International Journal of Game Theory*, 2: pp. 65-67.
- Schram, Arthur and Joep Sonnemans (1996a) "Voter Turnout as a Participation Game: An Experimental Investigation," *International Journal of Game Theory*, 225(3): pp. 385-406.
- Schram, Arthur and Joep Sonnemans (1996b) "Why People Vote: Experimental Evidence," *Journal of Economic Psychology*, 17: pp. 417-442.
- Sundali, James A., Amnon Rapoport, and Darryl A. Seale (1995) "Coordination in Market Entry Games with Symmetric Players" *Organizational Behavior and Human Decision Processes*, 64): pp. 203-218.
- Tullock, Gordon (1967) "The Welfare Costs of Tariffs, Monopolies, and Thefts," *Western Economic Journal*, 5: pp. 224-232.
- Tullock, Gordon (1980) "Efficient Rent Seeking," in J. Buchanan, *et al.* eds., *Toward a Theory of the Rent Seeking Society*, College Station: Texas A&M Press.
- Van Huyck, J.B., R.C. Battalio, and R.O. Beil (1990) "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure" *American Economic Review*, 80: pp. 234-248.