

DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

THE EFFECT OF CANDIDATE QUALITY ON ELECTORAL EQUILIBRIUM: AN EXPERIMENTAL STUDY

Enriqueta Aragonés
Institut d'Anàlisi, Campus Universitat Autònoma

Thomas R. Palfrey
California Institute of Technology



SOCIAL SCIENCE WORKING PAPER 1138

June 2002

The Effect of Candidate Quality on Electoral Equilibrium: An Experimental Study

Enriqueta Aragonés

Thomas R. Palfrey

Abstract

When two candidates of different quality compete in a one dimensional policy space, the equilibrium outcomes are asymmetric and do not correspond to the median. There are three main effects. First, the better candidate adopts more centrist policies than the worse candidate. Second, the equilibrium is statistical, in the sense that it predicts a probability distribution of outcomes rather than a single degenerate outcome. Third, the equilibrium varies systematically with the level of uncertainty about the location of the median voter. We test these three predictions using laboratory experiments, and find strong support for all three. We also observe some biases and show that they can be explained by quantal response equilibrium.

JEL classification numbers: C72, C92, D72

Key words: candidate quality; experiments; spatial competition; quantal response equilibrium

The Effect of Candidate Quality on Electoral Equilibrium: An Experimental Study*

Enriqueta Aragonés

Thomas R. Palfrey

1 Introduction

Candidate quality is widely considered to be a critical variable in electoral competition.¹ It affects the decisions of politicians to run for office, campaign fundraising and advertising, voter behavior, election outcomes, and, ultimately, policy outcomes. While direct measurement of candidate quality is often elusive, few doubt its importance in electoral politics. Quality differences between two candidates can arise for many reasons, including charisma, office-holding experience, incumbency, advertising, scandal, and other non-policy dimensions that can affect the relative attractiveness of two candidates. Political scientists have demonstrated over several decades of careful empirical research the importance of these and other image factors, or the “valence dimension,” as it is referred to in numerous articles and books.²

It is obvious that, all else constant, high quality candidates will fare better than low quality candidates. What is less obvious, but equally important, is that quality differences produce significant changes in the nature of political competition. Recent papers by Ansolobehere and Snyder (2000), Aragonés and Palfrey (2002), and Groseclose (2001) report a number of theoretical results about the equilibrium properties of spatial competition between two candidates who differ in quality.³ The results are striking,

*Aragonés acknowledges financial support from the Spanish Ministry of Science and Technology, grant number SEC2000-1186. Palfrey acknowledges financial support from the National Science Foundation, grant number SES-0079301, and from the Hacker Social Science Experimental Laboratory at Caltech. The paper has also benefited from suggestions by audiences at Caltech, Columbia University, Princeton University, ITAM, GREQAM, the 2001 meeting of the American Political Science Association, University of Malaga, and Universitat Pompeu Fabra.

¹For example, most studies of the incumbency advantage in congressional elections identify challenger quality as a critical factor. See Bond, Covington, and Fleischer (1985), Green and Krasno (1988), Jacobson and Kernell (1981), Krasno (1994), Squire (1992), Maisel, Stone and Maestas (1999), and the references they cite. Incumbency itself can also be viewed as an indicator of quality.

²See, for example, Stokes (1963), Kiewiet (1983), Banks and Kiewiet (1989), Kiewiet and Zheng (1993), Popkin (1976), and the references they cite.

³There are also some earlier theoretical papers that studied related kinds of asymmetry, such as incumbency or partisanship. These include Adams (1998), Bernhardt and Ingberman (1985), and Londregan and Romer (1993).

and suggest that the indirect equilibrium effects of candidate quality differentials may be even more important in determining candidate policies and election outcomes than the direct effects of producing more votes for one candidate than the other.⁴

The main insight about spatial competition if the candidates differ in quality (or along some other valence dimension) is that *the better candidate has an incentive to copy the policies of the inferior candidate, or at least move in that direction, while the disadvantaged candidate has the opposite incentive, to distance himself from the advantaged candidate.* Theoretically, the advantaged candidate will win all the votes if the two candidates choose sufficiently similar policies. Thus, in the standard Downsian model where candidates are purely office-motivated, the disadvantaged candidate must mix in order not to be predictable. However, in order for mixing to be an equilibrium strategy for the disadvantaged candidate, the advantaged candidate also must be mixing.⁵

This implies the first of three key properties of equilibrium in these models: the equilibrium makes statistical predictions, not point predictions. If both candidates have complete information and symmetric beliefs about voters, then the equilibrium is generally in mixed strategies. If candidates have private information with continuous types, then this mixed equilibrium can be “purified.” That is, there will exist an equilibrium in pure strategies, where the equilibrium locations of candidates will vary with this private information. Moreover, both the pure and mixed strategies produce distributions of location decisions that share similar statistical properties (Aragones and Palfrey 2001).

The second key property is that the distribution of location decisions of the two candidates will be different from each other, and the differences are systematic. We call this “*the quality divergence hypothesis.*” The main difference between the two candidate locations is that the distribution of locations of the better candidate is concentrated in the center of the policy space (i.e. the expected location of the median voter), while the location of the disadvantaged candidate tends toward the extremes. That is, *better candidates tend to choose more moderate locations.* Groseclose (2001) notes that this is consistent with two regularities that have been identified in empirical studies of congressional elections. One is the lack of support for the marginality hypothesis, documented in Fiorina (1973). That is, Fiorina finds that incumbents who are in marginal districts tend to moderate *less* than incumbents from safe districts. This is clearly consistent with the quality divergence hypothesis. Second, there is a recent paper by Ansolobehere, Snyder, and Stewart (2001) who compare the spatial locations of three categories of candidates: (1) incumbents seeking re-election; (2) candidates for open seats; and (3) challengers

⁴Using a different approach, Banks and Kiewiet (1989) show that candidate quality differentials can have important and surprising equilibrium effects on the entry of challengers in Congressional elections.

⁵An equilibrium in mixed strategies is guaranteed to exist (Aragones and Palfrey, 2002), and has intuitive properties. A pure strategy equilibrium may exist if candidates obtain utility from winning policies as well as from holding office, under certain conditions. Groseclose(2001) studies the properties of stable pure strategy equilibria, under the maintained assumption that they exist, but does not characterize conditions for existence. He presents an example suggesting that existence is especially problematic for small-to-intermediate values of the quality advantage, and if office holding is the primary motivation of candidates. The properties of pure strategy equilibria are similar to those of mixed equilibria.

running against an incumbent. They find that incumbents are the most moderate, followed by open seat candidates, and that challengers adopt the most extreme positions. To the extent that quality correlates across these three categories as expected, then this is consistent with the we think of incumbents from safe districts as being advantaged due to unbalanced partisanship of the district, then this provides further corroboration of the quality divergence hypothesis.

The third property is that the two candidates' equilibrium distributions of locations varies systematically with the degree of uncertainty about the median voter. Uncertainty helps the disadvantaged candidate and leads to a less centrist location of the advantaged candidate. On average, reduced uncertainty (or less polarization of the electorate) will lead to more centrist outcomes. Put another way, the quality divergence effect is strongest when there is a lot of uncertainty or if the electorate is highly polarized, and the effect is weakest when there is little uncertainty or a high degree of consensus in the electorate. We call this "*the polarization hypothesis.*"

Because the nature of equilibrium is very subtle in these asymmetric spatial competition games, and because the equilibrium (with complete information) is mixed, one cannot help but be skeptical about whether the features of the theoretical equilibrium might actually occur in practice. While the evidence put forth by Groseclose (1973) is consistent with the quality divergence hypothesis, that evidence could be explained by other theories. For example, the correlation between incumbency (i.e. electoral success) and moderation is also consistent with the standard Downsian model, or the more general models by Calvert (1985) and Wittman (1983) that include policy motivations. Thus, the evidence is suggestive that the theory may be on the right track, but does not provide a conclusive test of the model. Unfortunately, the kind of field data one would need to test these predictions are simply not available, due to the difficulty of obtaining reliable and accurate measurement of the "quality" variable, the degree of uncertainty or polarization in the electorate, and the location of candidates, and because the statistical nature of predictions would require a large number of observations. We believe that direct testing of the theory is needed.

With this in mind, we designed and conducted laboratory experiments to directly test both the quality divergence hypothesis and the polarization hypothesis. By doing so, we hope to find out if the basic predictions of the theory are accurate, and, if not, what sort of modification of the theory might be required. This paper reports and analyzes the data from those experiments. There are two main findings. First, all of the main qualitative properties of the equilibrium were clearly observed in the data. Both the quality divergence hypothesis and the polarization hypothesis are strongly supported by the data. Location decisions were statistical; the advantaged candidates located more centrally on average; and all of the comparative static predictions of changes in the distribution of voters were observed. In particular, when the distribution of voters was more polarized, there was more divergence. Second, while the main hypothesis were clearly supported, the exact distribution of locations was somewhat different from the quantitative predictions of the theory in systematic and surprising ways. There were two

interesting differences. First, both candidates showed a bias toward centrist locations, relative to the theoretical predictions. Second, the disadvantaged candidate’s location distribution was less responsive to changes in uncertainty than predicted by the theory.

In order to account for these anomalies, we consider an extension of Nash equilibrium theory that allows for a limited amount of bounded rationality. This approach, called *Quantal Reponse Equilibrium* (McKelvey and Palfrey, 1995, 1996), is based on two principles. The first principle is that players of any game respond continuously, but imperfectly, to the incentive structure in the game. While they cannot always optimize perfectly, they will *on average* choose better strategies more often than worse strategies. The second principle is that players are aware that othe players are also imperfect, and take this into account in choosing their actions. This boundedly rational version of Nash equilibrium often leads to surprising and unintuitive predictions about behavior in games, and provides a statistical model for data analysis.⁶ To analyze our data using this approach, we fit the data to the Logit version of Quantal Response Equilibrium. We find that the simplest one-parameter version of that model provides an excellent fit to the data, and accounts for the two unexpected findings.

In section 2 we summarize the results of Aragones and Palfrey (2002) and present the specific model that we use for the laboratory study and solve for the Nash equilibrium. In section 3 we describe the details of the experimental design. In section 4, we present the results of the experiments. In section 5 we describe and estimate the QRE model using the experimental data. We conclude in section 6.

2 The Model

We begin by describing the model and basic results of Aragones and Palfrey (2002). The policy space, φ , is one-dimensional and consists of the set of n points on the $[0, 1]$ interval, $x_i = \frac{i-1}{n-1}$, $i = 1, 2, \dots, n$. There are two candidates, A and D , who are referred to as the advantaged candidate and the disadvantaged candidate, respectively. Each candidate’s objective is to maximize his probability of winning the election. Each voter has a utility function, with two components, a policy component, and a candidate image component.⁷ The policy component is characterized by an ideal point in the policy space φ , with utility of alternatives in the policy space a strictly decreasing function of the Euclidean distance between the ideal point and the location of the policy, symmetric around the ideal point. We assume there exists a unique median location, denoted by x_m . Candidates do not know x_m , but share a common prior belief about it. This commonly shared belief is represented by a probability distribution over φ , and is denoted by a vector of probabilities, (ρ_1, \dots, ρ_n) , where $\rho_i \geq 0$, $i = 1, 2, \dots, n$ and $\rho_1 + \dots + \rho_n = 1$. We denote by m the median of the distribution ρ . The image component is captured by

⁶For example of applications of Quantal Response Equilibrium approach, see Goeree and Holt (2001), Guarnaschelli, McKelvey, and Palfrey (2000), and Signorino (1999).

⁷There could be either a finite number of voters or a continuum.

an additive constant to the utility a voter gets if A wins the election. That is, the utility that a voter with ideal point x_i obtains if A wins the election is $U_i(x_A) = \delta - |x_i - x_A|$ and his utility if candidate D wins is $U_i(x_D) = -|x_i - x_D|$, where candidates' policy positions are denoted by x_A and x_B , and the size of A 's advantage is $\delta \in (0, \frac{1}{n-1})$.⁸

The game takes place in two stages. In the first stage, candidates simultaneously choose positions in \wp . In the second stage, each voter votes for the candidate whose election would give him the higher utility. In case of indifference, a voter is assumed to vote for each candidate with probability equal to $1/2$.⁹

A pure strategy equilibrium is a pair of candidate locations, (x_A, x_D) such that both candidates are maximizing the probability of winning, given the choices of the other candidate. We denote by $\pi_A(x_A, x_D)$ and $\pi_D(x_A, x_D)$ the probability of winning for candidate A and for candidate D , respectively, as a function of (x_A, x_D) .¹⁰ A mixed strategy equilibrium is a pair of probability distributions over \wp , (σ^A, σ^D) , such that there is no mixed strategy for A that guarantees higher probability of winning than σ^A , given σ^D and there is no mixed strategy for D that guarantees higher probability of winning than σ^D , given σ^A .

2.1 Properties of Equilibrium

There are six main results, each of which we state without proof and provide a brief explanation of the result.

Result 1: If $n > 1$, then there does not exist a pure strategy equilibrium.

There is never an equilibrium in pure strategies.¹¹ The intuition is simple. If the disadvantaged candidate's location is perfectly predictable, the advantaged candidate can copy that strategy and win for sure. Therefore, (at least) the disadvantaged candidate must be mixing.

For the next set of results, we need two definitions.

Definition 1: A strategy for candidate j is *symmetric* if $\sigma_{ji} = \sigma_{j,n-i+1}$.

Definition 2: A strategy for candidate j has *no gaps* if there exist integers i, k such that $0 \leq i \leq k \leq n$ and $\sigma_{jt} > 0$ if and only if $i \leq t \leq k$.

⁸We refer to this as the “small advantage” case. See Aragones and Palfrey (2002) and Aragones and Palfrey (2001) for discussion and results for the large advantage case.

⁹The results do not depend on the tie-breaking rule.

¹⁰This model is equivalent to assuming candidates maximize expected vote, with a uniform distribution of voter ideal points.

¹¹Versions of this proposition, can be found in Groseclose [12] and Berger, Munger, and Potthoff [5]. Ansolobehere and Snyder [2] contains some related results.

Result 2: Generically, there exists a unique equilibrium in symmetric mixed strategies with no gaps.

The possibility that there are asymmetric equilibria, or equilibria with gaps is not ruled out.

Result 3: The distribution of D 's symmetric no-gap equilibrium strategy is U -shaped. That is, $\sigma_k^D \geq \sigma_{k+1}^D$ for $1 \leq k < \frac{n}{2}$ and $\sigma_k^D \leq \sigma_{k+1}^D$ for $\frac{n}{2} \leq k \leq n-1$.

Result 4: The distribution of A 's symmetric no-gap equilibrium strategy is single peaked. That is, $\sigma_k^A \leq \sigma_{k+1}^A$ for $1 \leq k < \frac{n}{2}$ and $\sigma_k^A \geq \sigma_{k+1}^A$ for $\frac{n}{2} \leq k \leq n-1$.

Result 5: Generically, the supports of the symmetric no-gap equilibrium strategies of A and D are the same.

Result 6: The probability that A wins is greater than the probability that D wins.

Summarizing, the main results are that there exists an essentially unique equilibrium in symmetric mixed strategies with no gaps. In this equilibrium, A is the more likely candidate to win. That is, A 's quality advantage leads to an electoral advantage. The supports of the equilibrium mixed strategies are the same, but otherwise the two distributions of the two candidates' are much different. The better candidate is more likely to locate in the center of the policy space than at the extremes, while the opposite is true for the lower quality candidate. This last property follows from results 3 and 4, and we call it "the quality divergence effect."

2.2 The equilibrium effects of uncertainty

The results described above were derived under the assumption that the location of the median voter is uniformly distributed on the policy space. In order to study the effect of changes in uncertainty, we need a model which can be solved for arbitrary distributions of the median voter. We do this by looking at a model where candidates can either locate in the center, on the left, or on the right. While this is a simplification that does not allow finer gradations of ideological placement of candidates, it captures the essence of the problem, since the main difference between the two candidates' strategies is the frequency that they locate in the center position.¹²

With three possible locations, the equilibrium has been fully solved and is given below. Denote the three possible locations, L, C or R , to denote Left, Center, and Right. where $L < C < R$. The probability the median voter is located at ideal point in L is denoted by α , similarly the probabilities she is located at ideal points C or R are denoted by β and γ respectively, with $\alpha + \beta + \gamma = 1$. Suppose that the utility

¹²The model can be solved for finer gradations of ideology, but little insight is gained. See Aragonés and Palfrey (2002) for a treatment of the $n = 4$ case.

functions of the voters are as the described in the previous section, and assume that $|(R - C) - (C - L)| < \delta < \max\{(C - L), (R - C)\}$.¹³ To maintain symmetry in the problem, we assume that $\alpha = \gamma \leq \frac{1}{2}$. Since $\alpha + \beta + \gamma = 1$, this implies that $\beta = 1 - 2\alpha$, so the model is reduced to a single parameter, α , which is proportional to the variance of the distribution. Thus, we call α the *uncertainty* (or *polarization*) *index*. When $\alpha > \frac{1}{3}$, the distribution of the median voter's ideal point is bimodal. We refer to this as the case of *high* uncertainty. When $\alpha < \frac{1}{3}$, the distribution of the median voter's ideal point is unimodal. We refer to this as the case of *low* uncertainty. The case of $\alpha = \frac{1}{3}$ is called the *uniform* case.

Formally our model is equivalent to a model in which the population distribution of the voter ideal points is given by (α, β, γ) and each candidate maximizes expect vote. In this context, α can also be interpreted as a measure of the polarization of the electorate. If $\alpha > \frac{1}{3}$, the distribution of voters is bimodal, while if $\alpha < \frac{1}{3}$, the distribution of voters is more concentrated in the center. Because the two versions of the model yield results that are formally equivalent, we will use the terms “uncertainty” and “polarization” interchangeably, and both terms will refer to the level of α .

Under these assumptions, the payoff matrix for the game is given in Table 1 below:

	a	b	c
a	$1, 0$	$\alpha, 1 - \alpha$	$1 - \alpha, \alpha$
b	$1 - \alpha, \alpha$	$1, 0$	$1 - \alpha, \alpha$
c	$1 - \alpha, \alpha$	$\alpha, 1 - \alpha$	$1, 0$

Table 1. Payoff Matrix for 3×3 Game

where A is the row player and D is the column player.

The (unique) mixed equilibrium is solved in the standard way. For any mixed strategy by D , denoted σ^D , A 's expected payoffs for the three pure strategies are given by:

$$\begin{aligned}\pi^A(L|\sigma^D) &= \sigma_L^D + \alpha\sigma_C^D + (1 - \alpha)\sigma_R^D \\ \pi^A(C|\sigma^D) &= (1 - \alpha)\sigma_L^D + \sigma_C^D + (1 - \alpha)\sigma_R^D \\ \pi^A(R|\sigma^D) &= (1 - \alpha)\sigma_L^D + (\alpha)\sigma_C^D + \sigma_R^D\end{aligned}$$

To solve for the mixed strategy equilibrium we equate these expected payoffs:

$$\begin{aligned}\pi^A(L, \sigma^D) = \pi^A(C, \sigma^D) &\implies \alpha\sigma_L^D - (1 - \alpha)\sigma_C^D = 0 \\ \pi^A(C, \sigma^D) = \pi^A(R, \sigma^D) &\implies (1 - \alpha)\sigma_C^D - \alpha\sigma_R^D = 0\end{aligned}$$

¹³That is, the quality advantage is large enough, so a C -location voter will vote for L , when the two candidates choose opposite extremes, but small enough that voters at D 's location will vote for D unless A is also located there. If the quality advantage is outside this range, the equilibria are trivial and uninteresting.

Thus the equilibrium value of σ^D can be determined by solving the following system of 3 equations:

$$\begin{aligned}\alpha\sigma_L^D - (1 - \alpha)\sigma_C^D &= 0 \\ (1 - \alpha)\sigma_C^D - \alpha\sigma_R^D &= 0 \\ \sigma_L^D + \sigma_C^D + \sigma_R^D &= 1\end{aligned}$$

Following a similar logic, the equilibrium value of σ^A can be determined by solving a similar system of 3 equations. This yields the following pair of equilibrium strategies

$$\begin{aligned}\sigma_L^A &= \sigma_R^A = \frac{\alpha}{2 - \alpha}, \sigma_C^A = \frac{2 - 3\alpha}{2 - \alpha} \\ \sigma_L^D &= \sigma_R^D = \frac{1 - \alpha}{2 - \alpha}, \sigma_C^D = \frac{\alpha}{2 - \alpha}\end{aligned}$$

To simplify notation, we denote the equilibrium by (p^*, q^*) where $p^* = \sigma_C^A$ and $q^* = \sigma_C^D$, are the equilibrium probabilities that A and D locate in the center position, respectively. The probabilities of locating at L (or R) are therefore $\frac{1-p^*}{2}$ and $\frac{1-q^*}{2}$, respectively. Using this notation, the equilibrium is $p^* = \frac{2-2\alpha}{2-\alpha}$ and $q^* = \frac{\alpha}{2-\alpha}$.

This equilibrium solution has several interesting properties. First note that since $\alpha \leq \frac{1}{2}$ this implies that $\sigma_L^A \leq \sigma_C^A$ and $\sigma_L^D \geq \sigma_C^D$. This is the quality divergence effect: the advantaged candidate places more weight in the central location, and the opposite is true for the disadvantaged candidate. This was proved only for the uniform distribution ($\alpha = 1/3$) in the earlier section, but holds more generally in this 3-location model.

Second, the comparative statics with respect to the uncertainty index, α , are interesting and a bit surprising. First, $\frac{\partial p^*}{\partial \alpha} < 0$ so that as uncertainty increases (or the electorate becomes more dispersed) the advantaged candidate moves away from the central location. Surprisingly, the opposite is true for D . That is, $\frac{\partial q^*}{\partial \alpha} > 0$, implying that the disadvantaged candidate moves toward the center as uncertainty increases. At the extreme case, when $\alpha = \frac{1}{2}$ (i.e. zero probability that the median is in the center), both candidates mix uniformly over the three locations. In other words, *both candidates tend to moderate as the polarization index increases*. We call this the “polarization hypothesis.”

Finally, uncertainty benefits the weaker candidate. The equilibrium probability that D wins is given by:

$$\Pi_D^*(\alpha) = \frac{\alpha(1 - \alpha)}{2 - \alpha}$$

The change in this equilibrium probability as α changes is found by computing the derivative of $\Pi_D^*(\alpha)$, which is given by:

$$\frac{d\Pi_D^*(\alpha)}{d\alpha} = \frac{1 - 2\alpha}{2 - \alpha} + \frac{\alpha(1 - \alpha)}{(2 - \alpha)^2} > 0$$

The derivative is positive because $\alpha \leq \frac{1}{2}$.

3 Experimental Design and Implementation

We conducted laboratory experiments using 3 different values of α , corresponding to three different levels of uncertainty. The three values were $\alpha = 1/3$ (uniform), $\alpha = 1/5$ (low uncertainty), and $\alpha = 3/7$ (high uncertainty). The experiments used students from California Institute of Technology (CIT) and Universitat Pompeu Fabra (UPF). Nine sessions were conducted, three for each value of α , of which two were carried out at CIT and one at UPF. Table 2 summarizes the information about each session.

<i>Uncertainty</i>	(p^*, q^*)	<i># subjects</i>	<i>UPF</i>	<i>CIT</i>	<i># rounds</i>
Uniform ($\alpha = \frac{1}{3}$)	$(\frac{3}{5}, \frac{1}{5})$	32			
Session 1		10		x	200
Session 2		8		x	200
Session 3		16	x		200
Low ($\alpha = \frac{1}{5}$)	$(\frac{7}{9}, \frac{1}{9})$	40			
Session 1		14		x	200
Session 2		10		x	200
Session 3		16	x		200
High ($\alpha = \frac{3}{7}$)	$(\frac{5}{11}, \frac{3}{11})$	38			
Session 1		8		x	166
Session 2		14		x	200
Session 3		16	x		200

Table 2. Experiment session summary

3.0.1 Procedures

The experiments were conducted using the PLDK software developed at the Hacker Social Science Experimental Laboratory at Caltech. The interface for the software presents each subject with a matrix of payoffs, and keeps track of the history of previous game outcomes automatically for each subject. The matrices of payoff were games that are strategically equivalent to the 3-location games, but constants were added to the payoffs to avoid zero outcomes and to approximately equalize the expected earnings for A and D players. See Appendix B for the matrix of payoffs.

Each session lasted 200 rounds, each round being one play of the game.¹⁴ Between 8 and 16 students participated in each session. Total earnings were equal to the sum of all earnings over the 200 rounds. Average earnings ranged from \$20 to \$25 and sessions lasted about 90 minutes.¹⁵

Subjects played both roles (A and D). At the beginning of the session, subjects were assigned roles as either row or column players and instructions were read aloud. The

¹⁴The first High session was terminated after round 166 due to a computer crash.

¹⁵For the experiments conducted at UPF, earnings averaged 13 Euros.

game matrix was displayed in front of the room for everyone to see. It also appeared on their computer screen. In each round row players clicked their mouse on a row to make a decision and column players clicked on a column to select their decision. After everyone had made a decision, the row/column outcome of their match was highlighted in the matrix on their screen. The screen also kept a display of the history of their play and the choices made by their past opponents. Several practice rounds were conducted in order to familiarize the subjects with the computer interface. During these practice rounds, the subjects were not allowed to make any choices on their own.

The subjects then played 100 rounds, being randomly rematched into pairs (one column player and one row player) after each round of play. After round 100, the payoff matrix was changed so that the row and column players payoffs were reversed. This reversal was carefully explained to the subjects. They played 100 additional times with these reversed payoffs. This reversal allowed each subject to have 100 rounds of experience as the A player and 100 rounds of experience as the D player. Subjects were told all of this information in advance. The instructions were worded in neutral terms that would not be associated with personal political ideology. The three strategies were labeled “A, B, C”. A sample copy of the instructions is attached as an appendix to this paper.

3.0.2 Hypotheses

We have the following comparative static hypotheses that are derived from the theory. These are summarized below, denoting the empirical choice frequencies of center by (\hat{p}, \hat{q}) , and the treatments by H, M, L (for High, Medium, or Low uncertainty)

1. $\hat{p}_L > \hat{p}_M > \hat{p}_H$
2. $\hat{q}_L < \hat{q}_M < \hat{q}_H$
3. $\hat{q} < \frac{1}{3} < \hat{p}$ for all uncertainty treatments

These hypotheses can be summarized in one string of inequalities: $\hat{q}_L < \hat{q}_M < \hat{q}_H < \frac{1}{3} < \hat{p}_H < \hat{p}_M < \hat{p}_L$

4 Experimental Results

4.1 Effect of uncertainty on location decisions

The main results of the paper have to do with the effect of the primary treatment variable (level of uncertainty) on the location decisions of the subject candidates. Recall that there are four primary hypotheses based on the Nash equilibrium model presented in the previous section. The first hypothesis states that disadvantaged candidates will choose

the center location *less* than one-third of the time, *regardless of the level of uncertainty*. The second hypothesis states that advantaged candidates will choose the center location *more* than one-third of the time, *regardless of the level of uncertainty*. Thus hypotheses 1 and 2 jointly imply the **quality divergence hypothesis**. The third hypothesis states that disadvantaged candidates will choose the center location *more often as the level of uncertainty increases*. The fourth hypothesis states that advantaged candidates will choose the center location *less often as the level of uncertainty increases*. Thus hypotheses 3 and 4 jointly imply the **polarization hypothesis**. Table 3 compares the theoretically predicted center choice probabilities the aggregate relative choice frequencies by subjects when they were candidates *A* and *D* for each of the three treatments.

<i>Uncertainty</i>	p^*	q^*	\hat{p}	\hat{q}
Uniform <i>N</i>	.6	.2	.609 3200	.288 3200
Low <i>N</i>	.78	.11	.769 4000	.252 4000
High <i>N</i>	.45	.27	.514 3664	.320 3664

Table 3. Nash predictions and Aggregate Date

The key observation from this table is that all four of the main hypotheses of the theory are strongly supported by the aggregate data from this experiment. The order of center choice relative frequencies, $\hat{q}_L < \hat{q}_M < \hat{q}_H < \frac{1}{3} < \hat{p}_H < \hat{p}_M < \hat{p}_L$, is exactly what is predicted by Nash equilibrium. Table 4 clearly shows the support for all of these theoretical hypotheses. That table displays the differences between pairs of aggregate choice frequencies. The cell entries in the table correspond to the difference is the relative frequency of center choices for two treatments (or one treatment compared with $\frac{1}{3}$). For example, the entry in the cell with row label \hat{p}_L and column label \hat{p}_M is $\hat{p}_L - \hat{p}_M = .769 - .609 = .160$. Every single one of the 21 predicted differences have the correct sign, and all except one are statistically significant at better than the 1% level.

	\hat{p}_L	\hat{p}_M	\hat{p}_H	$\frac{1}{3}$	\hat{q}_H	\hat{q}_M	\hat{q}_L
\hat{p}_L		.160**	.255**	.436**	.449**	.481**	.517**
\hat{p}_M			.095**	.276**	.289**	.321**	.357**
\hat{p}_H			—	.181**	.194**	.226**	.262**
$\frac{1}{3}$				—	.013*	.045**	.081**
\hat{q}_H					—	.032**	.068**
\hat{q}_M						—	.036**
\hat{q}_L							—

Table 4. The entries in each cell correspond to the difference between row and column in relative frequency of choosing Center.
 ** significant at better than 1% level.
 * significant at better than 5% level.

With the exception of the bimodal treatment, the aggregate fit for the A players to the quantitative prediction of Nash equilibrium was nearly perfect. The quantitative fit for the D players is not nearly as good, and the error was in the direction of overplaying the center strategy in all cases. The A players overplayed the center strategy in 2 of 3 treatments (the exception being the *High* treatment, where the difference is very small.) In addition, the D players do not respond very strongly to the treatment effects. That is, the differences between \hat{q} in the different treatments was always less than predicted by the theory. Thus, while the qualitative features of the data are very supportive of the theory, the actual magnitudes of (\hat{p}, \hat{q}) in the various treatments deviate somewhat from the Nash equilibrium predictions in systematic ways. This can be seen clearly in Figure 1, which displays the Nash predictions and the relative frequencies from the aggregate data, by treatment. In that figure, p is on the horizontal axis and q is on the vertical axis.

FIGURE 1 ABOUT HERE

To summarize, there are five main features of the aggregate data:

1. The quality divergence hypothesis is strongly supported by the data.
2. The polarization hypothesis is strongly supported by the data.
3. All of the signed comparative static predictions about p and q were strongly supported by the data.
4. All of these comparative static differences are statistically significant.
5. The A player fits the Nash predictions much better than the D player.
6. Both players tend to overplay the center strategy, and this effect is strongest for the D players.

7. The response of the D players to changes in the level of uncertainty is less than predicted.

4.2 Quantal Response Equilibrium Analysis

The strategic structure of equilibrium suggests the following possible explanation. If D players begin with uninformative prior beliefs about the choices by A players, then locating in the center is their *optimal* choice. The same is true for the A players. This could produce a pattern in which both players initially overplay C , and then gradually adapt in the direction of their equilibrium strategy. Since in this kind of process D starts out further away from his equilibrium strategy than A , it is not surprising that A frequencies are closer to their equilibrium values than the D frequencies. What is needed to capture this idea theoretically is a model that can predict one player to be further from Nash equilibrium than the other player.

Quantal Response Equilibrium (QRE) is an equilibrium model of imperfect play. A quantal response function is simply a smoothed out single-valued best response function that is monotonically increasing in expected payoffs. The quantal response functions are continuous and “statistical” in the sense that all strategies are played with positive probability. Therefore, players do not *always* play best responses. However, the monotone property implies that they play better strategies more frequently than worse strategies. Formally, for each player, the quantal response function maps the vector of expected payoffs of feasible actions into mixed strategy, that satisfies monotonicity and continuity properties. A quantal response equilibrium is a fixed point of the following composed mapping. Let σ be some (mixed) strategy profile in the game. Given σ , one can compute, for each player i and for each of player i 's possible actions j , the expected payoff from playing that action, given σ , denoted U_{ij} . Given these vectors of expected payoffs, the quantal response functions of players then yield a new mixed strategy profile, $\hat{\sigma} = QR(\sigma)$. A QRE is a fixed point of this mapping, that is, a mixed strategy profile, σ^* , with the property that $\sigma^* = QR(\sigma^*)$. McKelvey and Palfrey (1995) establish a number of theoretical properties of QRE points, including existence and upper hemicontinuity and a connection between QRE and Bayesian equilibrium of games with payoff disturbances.

A particularly useful parametric form of QRE is the Logit equilibrium. The Logit equilibrium arises when all players' quantal response functions are Logit functions of the expected utilities that are implied by the mixed strategies. Formally, a Logit quantal response function is given by:

$$\hat{\sigma}_{ij} = \frac{e^{\lambda U_{ij}(\sigma)}}{\sum_k e^{\lambda U_{ik}(\sigma)}}$$

where λ is a parameter measuring the responsiveness of i to payoff differences between strategies. A Logit equilibrium is therefore a mixed strategy profile σ^* such that

$$\sigma_{ij}^* = \frac{e^{\lambda U_{ij}(\sigma^*)}}{\sum_k e^{\lambda U_{ik}(\sigma^*)}} \text{ for all } i \text{ and } j$$

When $\lambda = 0$, behavior is completely unpredictable, and the unique Logit equilibrium has every player choosing actions according to a uniform distribution. When $\lambda \rightarrow \infty$, the Logit equilibria converge to Nash equilibria. The Logit equilibrium correspondence for a game is the set of all Logit equilibria for the game, for all non-negative values of λ . Because of its relatively simple functional form, Logit equilibria are relatively easy to compute numerically, and in some cases analytically.

Figures 2, 3, and 4 graph the Logit equilibrium correspondences for the H, M, and L treatments, respectively. Choice probabilities are on the vertical axis and λ is on the horizontal axis. For simplicity, we only include the choice frequencies for *C* locations, one for each of the two players. Thus, each graph has two curves, one for the *A* player and one for the *D* player. From the graphs, one can easily see how the QRE captures the intuition that the *A* players converge more quickly to the Nash equilibrium, while the *D* players converge slowly. Neither converges monotonically. For intermediate values of λ , both players overplay *C* relative to Nash equilibrium.

FIGURE 2 ABOUT HERE
 FIGURE 3 ABOUT HERE
 FIGURE 4 ABOUT HERE

The Logit equilibrium correspondence also provides a structural model that permits us to fit the data to QRE using standard maximum likelihood techniques. Given a dataset consisting of n observations of *A* and *D* choices in the location game, one can construct the likelihood function as a function of the free parameter, λ , which is determined by the theoretical choice probabilities of the unique Logit equilibrium for that value of λ . The maximum likelihood estimate of λ is the value of λ at which that likelihood function is maximized. This parameter estimate in turn implies estimated equilibrium choice frequencies, (\hat{p}^*, \hat{q}^*) using the formula above.

Figure 5 is similar to figure 1, but also includes the fitted QRE-predicted choice probabilities of *A* and *D*, in addition to the Nash predictions and the aggregate data. The QRE model clearly picks up the three anomalous features of the data: overplay of *Center* by both players, and the worse fit of *D* compared to *A*, and the weaker responsiveness by *D* to changes in uncertainty.

FIGURE 5 ABOUT HERE

Table 5 presents the QRE estimates of the data broken down by treatment and model. Column one lists the three treatments, Uniform, Low, and High. We computed estimates for three different models, which we call the unconstrained model, the constrained model, and the Nash model, respectively. The unconstrained estimates allow a separate estimate of λ for each treatment, while the constrained estimate forces $\hat{\lambda}$ to be the same for all treatments. The Nash model computes the likelihood function using the Nash equilibrium

<i>Treatment</i>	<i>Model</i>	\hat{p}^*	\hat{q}^*	\bar{p}	\bar{q}	$\hat{\lambda}$	$-\log L$	χ^2 statistic
Uniform	unconstrained	.616	.289	.609	.288	1.04	6510	
	constrained	.618	.266			1.43	6515	8.70*
	Nash model	.600	.200			∞	6580	140*
Low Uncertainty	unconstrained	.803	.261	.769	.252	1.55	7150	
	constrained	.800	.272			1.43	7151	2.66+
	Nash model	.778	.111			∞	7441	582*
High Uncertainty	unconstrained	.462	.312	.514	.320	1.22	7817	
	constrained	.462	.306			1.43	7817	0.92+
	Nash model	.455	.273			∞	7839	44*

Table 5. QRE estimation results and tests of constrained models.
* Model restriction rejected at 1% significance level or better.
+ Model restriction cannot be rejected at 10% significance level.

choice probabilities, which correspond to the limit of the QRE choice probabilities when $\lambda \rightarrow \infty$.

Columns 3 and 4 give (\hat{p}^*, \hat{q}^*) , the estimated choice probabilities under the various models. Columns 5 and 6 give (\bar{p}, \bar{q}) , the empirical relative frequencies observed in the experiment. Column 7 gives the maximum likelihood estimate of λ , and column 8 gives minus the value of the log likelihood function for the model, evaluated at the maximum likelihood estimate of λ .

The constrained model is nested in the unconstrained model, so we use a likelihood ratio test to test for model rejection. The chi-square statistic (twice the log of the likelihood ratio), is given in the last column of table 5. While the unconstrained model fits slightly better in all three cases, the improvement in fit is insignificant (at the 10% level) for two of the treatments (Bimodal and Unimodal). Only with the Uniform treatment it is statistically significant (at the 1% level), but even for this case, the improvement in fit is of little real consequence, as the implied differences for choice probabilities between the two models are negligible.

The Nash equilibrium model is also (approximately) nested in the unconstrained model, so we can again use a likelihood ratio test to test for model rejection. The Nash model is easily rejected for all treatments, and the differences are statistically significant at any conventional level.

We next turn to the secondary hypotheses of the experiment, concerning *experience* and *heterogeneity*.

4.3 Experience and Learning

We investigate learning at a macro scale, simply asking whether aggregate behavior was different after subjects had a chance to observe the pattern of behavior of their opponents. Recall that our design used random matching, so that subjects were not trying to outguess an opponent based on observation of that opponent's play. Instead, the subjects were receiving information about the average play of the population over time. For this reason,

we focus our attention in this subsection on changes of average play over the course of a session.

We divide the data into two data subsets, one which we call *experienced* and the other which we call *inexperienced*. Since subjects played 100 rounds in each candidate's role, we define inexperienced rounds to be the first 50 rounds a subject was in a particular round and define experienced rounds to be the remaining 50 rounds a subject was in that role.¹⁶ There were clear and significant trends in the data, with the experienced data being closer to Nash equilibrium and also fitting the QRE model better. Also, in all cases the estimate of $\hat{\lambda}$ increases with experience, and the changes are significant at the 1% level or better. However, in most cases, the actual movement in the aggregate choice probabilities was not very large. Table 6 displays the estimates broken down separately by treatment and by experience level. Figure 6 displays the differences in a graph similar to figure 5.

		\hat{p}^*	\hat{q}^*	\bar{p}	\bar{q}	$\hat{\lambda}$	$-\log L$
Uniform	inexperienced	.611	.307	.623	.307	.856	3257
	experienced	.618	.270	.594	.269	1.35	3251
Low Uncertainty	inexperienced	.800	.270	.763	.260	1.46	3605
	experienced	.804	.251	.775	.245	1.68	3544
High Uncertainty	inexperienced	.461	.315	.516	.324	1.15	4051
	experienced*	.462	.310	.512	.315	1.32	3765

Table 6. QRE estimates (\hat{p}^*, \hat{q}^*) of the aggregate data (\bar{p}, \bar{q})
 *272 fewer observations than previous row

FIGURE 6 ABOUT HERE

4.4 Heterogeneity

This section examines variation in choice behavior across subjects. We find evidence for heterogeneity. Even with the heterogeneity, all of the comparative statics results are still supported in the data. This is important to note because the statistical tests in the earlier section assume that all observations are independent, and therefore the significance levels are inflated. One way to adjust for this is to conduct similar tests with the individual level data, comparing the population distribution of choice probabilities across samples, using non-parametric statistics. This is what we do here, with the pooled sample of individuals. Figure 7 shows the cumulative distributions of individual choice frequencies by treatment and by role. For example, in the uniform treatment, there were 32 subjects, so the graph shows 32 center choice frequencies for the individuals when they were A players and 32 choice frequencies for the same individuals when they were D players. Each point on the graph gives the relative frequency (out of 100 moves) that a particular individual chose the center strategy in a particular rule. The points are ordered by relative frequency

¹⁶In the case of the experiment that crashed, we lost 34 rounds of experienced data.

(not subject), so that the curves represent empirical cumulative distribution functions of individual choice frequencies. There is a clear ordering of these empirical CDF's, as hypothesized.

FIGURE 7 ABOUT HERE
FIGURE 8 ABOUT HERE

A second issue with heterogeneity arises because we used two separate subject pools. This feature of the design was implemented as a robustness check. The student populations (and culture) at UPF and CIT are different in many ways, but the theoretical model is intended to apply to both subject pools, so we do not predict a difference. Figure 8 displays the UPF and CIT data as well as the Nash predictions and the QRE estimates.

FIGURE 9 ABOUT HERE

Indeed the behavior in the two subject pools is very similar, in spite of some small quantitative differences. There is one reversal of the sign predictions, which occurs in the CIT data. That reversal is $\hat{q}_{L,CIT} > \hat{q}_{M,CIT}$, but the difference ($-.016$) is not significant at the 5% level.

5 Conclusions

The results of this laboratory experiment provide strong support for the theoretical equilibrium effects of candidate quality on policy location. The central predictions of the theory are the quality divergence hypothesis and the polarization hypothesis. That is: 1) both candidates diverge from the center, with the weaker candidate diverging more than the stronger candidate; and 2) as the distribution of voters becomes more spread out, both candidates moderate their positions. The design allowed us to test these key predictions about how endogenous variables (candidate locations) co-vary with candidate quality, and with the distribution of voters. All of these predictions were supported by the data. Altogether, the design and the model offer 21 predicted sign differences in the observable choice frequencies by the candidates, across the three treatments. *Every single one of these 21 predicted sign differences were the right sign and were statistically significant.*

The quantitative predictions of the theory were also the right order of magnitude, but there were two interesting biases that were observed. First, we found that when subjects were in the role of the advantaged candidate were more responsive to the changes in the distribution of voters than when they were in the role of the disadvantaged candidate. As a result the Nash equilibrium predictions fit the data for advantaged candidates better than the data for the disadvantaged candidates. Second both players tend to adopt

more moderate positions than was predicted by the model. We show that both of these observations can be accounted for very well by a bounded rationality version of Nash equilibrium, called *Quantal Response Equilibrium*.

6 References

1. Adams, James, “Policy Divergence in Multicandidate Probabilistic Spatial Voting,” *Public Choice* **100** (1998), 103–22.
2. Ansolabehere, Stephen and James M. Snyder, Jr., “Valence Politics and Equilibrium in Spatial Election Models,” *Public Choice* **103** (2000), 327–336.
3. Ansolabehere, Stephen, James M. Snyder, Jr. and Charles Stewart III, “Candidate Positioning and U. S. House Elections,” *American Journal of Political Science* **45** (2001), 136–59.
4. Aragonés, Enriqueta and Thomas R. Palfrey, “Mixed Equilibrium in a Downsian Model with a Favored Candidate,” *Journal of Economic Theory*, **103** (2002), 131–61.
5. Aragonés, Enriqueta and Thomas R. Palfrey, “Spatial Competition Between Two Candidates of Different Quality: The Effects of Candidate Ideology and Asymmetric Information,” Working Paper, Caltech, June 2001.
6. Banks, Jeffrey S. and D. Roderick Kiewiet, “Explaining Patterns of Candidate Competition in Congressional Elections,” *American Journal of Political Science* **33** (1989), 997–1015.
7. Berger, Mark. M., Michael C. Munger, and Richard F. Potthoff, “With Uncertainty, the Downsian Model Predicts Divergence,” *Journal of Theoretical Politics* **12** (2000), 262–268.
8. Bernhardt, Daniel and Daniel Ingberman, “Candidate Reputations and the Incumbency Effect,” *Journal of Public Economics* **27** (1985), 47–67.
9. Bond, Jon, Cary Carrington, and Richard Fleisher, “Explaining Challenger Quality in Congressional Elections,” *Journal of Politics* **47** (1985), 510–29.
10. Calvert, Randall, “Robustness of the Multidimensional Voting Model: Candidate Motivations, Uncertainty, and Convergence,” *American Journal of Political Science* **29** (1985), 69–95.
11. Downs, Anthony, *An Economic Theory of Democracy*, New York: Harper and Row, 1957.
12. GAMBIT software, <http://www.hss.caltech.edu/gambit>, 1999.
13. Goeree, Jacob and Charles Holt, “Ten Little Treasures of Game Theory and Ten Intuitive Contradictions” (with C. Holt), *American Economic Review*, **10** (2001), 1402–22.
14. Green, Donald and Jonathan Krasno, “Salvation for the Spendthrift Incumbent: Re-estimating the Effects of Campaign Spending in House Elections,” *American Journal of Political Science* **32** (1988), 884–907.

15. Groseclose, Tim, "A Model of Candidate Location When One Candidate Has a Valence Advantage" , *American Journal of Political Science* 45 (2001), 862-86.
16. Guarnaschelli, Serena, Richard D. McKelvey, and Thomas R. Palfrey, "An Experimental Study of Jury Decision Rules," *American Political Science Review* 94 (2000), 407-23.
17. Ingberman, Daniel, "Incumbent Reputations and Ideological Campaign Contributions in Spatial Competition," *Mathematical and Computer Modeling* 16 (1992), 147-169.
18. Jacobson, Gary and Samuel Kernell, *Strategy and Choice in Congressional Elections*, New Haven: Yale University Press, 1981.
19. Kiewiet, D. Roderick, *Macroeconomics and Micropolitics*, University of Chicago Press, Chicago, 1983.
20. Kiewiet, D. Roderick and Langche Zeng, "An Analysis of Congressional Career Decisions, 1947-86," *American Political Science Review* 87 (1993), 928-941.
21. Krasno, Jonathan, *Challengers, Competition, and Reelection*, New Haven: Yale University Press, 1994.
22. Londregan, John, and Thomas Romer, "Polarization, Incumbency, and the Personal Vote," in *Political Economy: Institutions, Competition, and Representation* (William Barnett, Melvin Hinich, and Norman Schofield, eds.), New York: Cambridge University Press, 1993.
23. Maisel, Sandy, Walter Stone, and Cherie Maestas, "Reassessing the Definition of Quality Candidates," *Paper prepared for Annual Meeting of the Midwest Political Science Association*, 1999.
24. McKelvey, Richard D. and Thomas R. Palfrey, "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 10 (1995), 6-38.
25. McKelvey, Richard D. and Thomas R. Palfrey, "A Statistical Theory of Equilibrium in Games," *Japanese Economic Review*, 47 (1996), 186-209.
26. Signorino, Curtis, "Strategic Interaction and Statistical Analysis of International Conflict," *American Political Science Review* 93 (1999), 279-98.
27. Squire, Pervill, "Challenger Quality and Voting Behavior in Senate Elections," *Legislative Studies Quarterly*, 17 (1992), 247-63.
28. Stokes, Donald E., "Spatial Models of Party Competition," *American Political Science Review* 57 (1963), 368-377.
29. Wittman, Donald, "Candidate Motivations: A Synthesis of Alternative Theories," *American Political Science Review* 77 (1983), 142-57.

7 Appendix: Sample Instructions

7.1 Introduction

Welcome to the SSEL Lab. Please do not do anything with the computer equipment until you are instructed to. Please put all of your personal belongings away, so we can have your complete attention. Raise your hand if you need a pencil. Feel free to adjust your chairs so they are comfortable for you.

This is an experiment in decision making, and you will be paid for your participation in cash, at the end of the experiment. Different subjects may earn different amounts. What you earn depends partly on your decisions and partly on the decisions of others.

The entire experiment will take place through computer terminals, and all interaction between you will take place through the computers. It is important that you not talk or in any way try to communicate with other subjects during the experiment.

We will start with a brief instruction period. During this instruction period, you will be given a complete description of the experiment and will be shown how to use the computers. If you have any questions during the instruction period, raise your hand and your question will be answered so everyone can hear.

7.2 Computer Instructions and Practice Matches

During the computer instruction session, we will teach you how to use the computer by going through a practice session. We will go through this practice session very slowly and it is important that you follow instructions exactly. *Do not hit any keys until you are told to do so, and when you are told to enter information, type exactly what you are told to type.* You are not paid for the practice session.

We will first pass out the practice experiment record sheet, on which you will record all of the results from this experiment. Please record your name, the date, and your social security number on the bottom of the sheet. Note that you have been assigned a color, either Red or Blue. The color is written on top of the record sheet.

[PASS OUT RECORD SHEETS]

[WAIT FOR SUBJECTS TO RECORD INFORMATION]

START PLDK SERVER PROGRAM ON SERVER, IF NOT DONE ALREADY]

Please click on the ICON that says "PLDK client." When the computer prompts you for your name, type your full name, your social security number, and click on your color. Then click OK to confirm. If you have any questions about how to do this, please raise your hand.

[WAIT FOR SUBJECTS TO LOG ON]

You now see the experiment screen. You have each been assigned to be either a RED subject or a BLUE subject in this experiment. Your color as well as your subject ID number is shown in the banner at the top of the screen. Please record your subject ID number on your record sheet.

[WAIT FOR SUBJECTS TO RECORD INFORMATION]

Each of you has been matched by the computer with a subject of the opposite color. If you are a BLUE subject, you are matched with one of the other RED subjects. If you are a RED subject, you are matched with one of the BLUE subject.

In the upper left part of the screen, you see a table.

[SHOW TABLE ON OVERHEAD PROJECTOR]

Will all subjects now move the mouse into the table and click it. If you are a RED subject, one of the rows will be highlighted. If you are a BLUE subject, one of the columns will be highlighted.

Each of you is asked to make a choice, but please do not do so at this time. If you are a RED subject, on the left of the screen you are asked to please choose a row. If you are a BLUE subject, you are asked to please choose a column. The outcome, and your payoff, is determined by the cell in the table that is chosen. In each cell of the table, the first number is the payoff for the RED subject, and the second number is the payoff to the BLUE subject.

[GO THROUGH A COUPLE CELLS IN OVERHEAD TABLE TO EXPLAIN]

[WAIT FOR SUBJECTS TO HIGHLIGHT ROW OR COLUMN]

Will all RED subjects now please choose “B” and all BLUE subjects please choose “A” by clicking the mouse button now while the arrow is pointing to the appropriate row or column. After you have made your choice, you are given a chance to confirm your decision. If it is not correct, please change it. When it is correct, please confirm by clicking on “confirm” now.

[WAIT FOR SUBJECTS TO CHOOSE AND CONFIRM CHOICE]

[WALK AROUND ROOM TO CHECK]

After all subjects have confirmed their choices, the match is over and you are shown the choice of the blue subject you were matched with. The outcome of the round, BA, is now highlighted in purple on everybody’s screen. Your earnings are determined by the entries in the highlighted cell of the table that was selected. So the payoff to a RED subject for the first match is 6 points and the payoff to a BLUE subject is 14 points. You are not being paid for the practice session, but if this were the real experiment, then the payoff you have recorded would be money you have earned from the first match, in points.

We will now proceed to the second practice match. Each match is the same except you are matched with a new subject of the opposite color. Note that the decisions and payoffs of the first match are recorded in the experiment history at the right side of the

screen. The outcomes of all of the previous matches will be recorded at the right side of the screen throughout the experiment so that you can refer back to previous rounds whenever you like.

[HIT KEY TO START SECOND ROUND]

For the second match, each of you have now been rematched with a new subject of the opposite color. All RED subjects again choose “B” and confirm. All BLUE subjects choose “C” by clicking on the right column.

[WAIT FOR SUBJECTS TO CHOOSE AND CONFIRM CHOICE]

The payoff to a RED subject for this practice round is 8 points and the payoff to a BLUE subject is 12 points. This concludes the second round. Notice that the results are again recorded in the history screen. Note also that the history screen keeps track of the number of times you have chosen each row or column, and of the average payoff you received from each row or column. For example, the red subject has chosen “A” twice. The first time she received 6, and the second time 8 points. So the average is 7.

[DO 4 MORE ROUNDS, CHOOSING (BA), (CC), (BB), (CA)]

[HIT KEY TO END PRACTICE SESSION]

This concludes the practice session. The computer screen now indicates your total points that you earned in the practice session. This is multiplied by the exchange rate to get your money Payoff. Since this is a practice session, the exchange rate is zero. In the actual experiment, the exchange rate is .01, so that each point is worth one cent.

[WAIT FOR SUBJECTS TO RECORD OUTCOME AND CLICK ”OK”]

7.3 Part 1

The actual experiment consists of two parts. Each part will last for 100 matches. When the first part is over, we will give you some additional instructions before the second part begins. Each match will proceed as in the first practice match, except you will be paid one cent for each point. The table will have three rows and three columns, the row and column labels will be the same as the practice, and the payoffs in the table will be the same as in the practice. Also, just like in the practice round, you will be randomly rematched with a new subject of the opposite color after each match.

The total amount you earn in this first part of the experiment is equal to the sum of your earnings in all 100 matches. You will be paid in cash at the end of the experiment. No other participant will be told how much money you earned in the experiment. You need not tell any other participants how much you earned. Are there any questions before we begin the experiment?

[TAKE QUESTIONS]

O.K., then we will now begin with the actual experiment.

Please lower your chairs to the lowest position, and pull out the dividers as far as they will go. This ensures your privacy and the privacy of the others in the experiment. We will now begin match number 1.

[START EXPERIMENT]
[AFTER FIRST MATCH ANNOUNCE:]

The first real match of the experiment is over. There will be 99 more matches in Part 1. Remember that you are randomly rematched with a new subject after every single match. After match 100 has finished, please record your total payoffs on your record sheet and then wait for the instructions for the second part of the experiment.

7.4 Part 2

This is the second and final part of the experiment. The amount of money you earn in this part will be added to the amount you earned in part 1 to determine your total money earnings for the whole experiment. Just as in part 1, each point is worth one cent. This part of the experiment is similar to the first part, except the payoff table has been changed in a very specific way.

[SHOW NEW TABLE ON OVERHEAD PROJECTOR NEXT TO OLD TABLE]

This is a payoff table that reverses the roles of Blue and Red. That is, Blue's payoffs are the same as Red's payoffs were in Part 1, and Red's payoffs are the same as Blue's payoffs were in Part 1. For example, suppose that Blue chooses A and Red chooses B. Then Blue gets 6 and Red gets 14. Now compare this to the payoffs in the first table, when Red chose A and Blue chose B.

[PUT UP COMBINED SLIDE WITH BOTH PAYOFFS]

Then Red got 6 and Blue got 14. If you look carefully at the new table, you will notice that it is derived from the old one by transposing it (that is, flipping it around the diagonal) and reversing the Red and Blue payoffs.

[ILLUSTRATE HOW THIS TRANSPOSITION WORKS USING OVERHEADS]

Are there any questions before we begin?

[TAKE QUESTIONS]

O.K., then we will now begin the second part. The second part will also have 100 matches. Remember that you are randomly rematched with a new subject after every single match. After match 100 has finished, please record your Part 2 total payoffs on your record sheet and then wait for instructions for how to be paid. We ask you to refrain from talking with each other, not only during the matches, but also while you are waiting to be paid. Thank you in advance for your cooperation.

[BEGIN SECOND PART]
[END AFTER MATCH 100]

The experiment is now over. Please record your part 2 money earnings and add them

to your part 1 money earnings. Enter this sum in the row labeled Total Earnings. You will be paid this amount of money in the next room. We will pay you one at a time, beginning with subject number 1. We ask you not to talk with each other or use the computer equipment while you are waiting to be paid. Subject number 1, will you please come with us to the next room. Please collect your belongings and bring them and your record sheet with you. You will be leaving from the outside door in the next room.

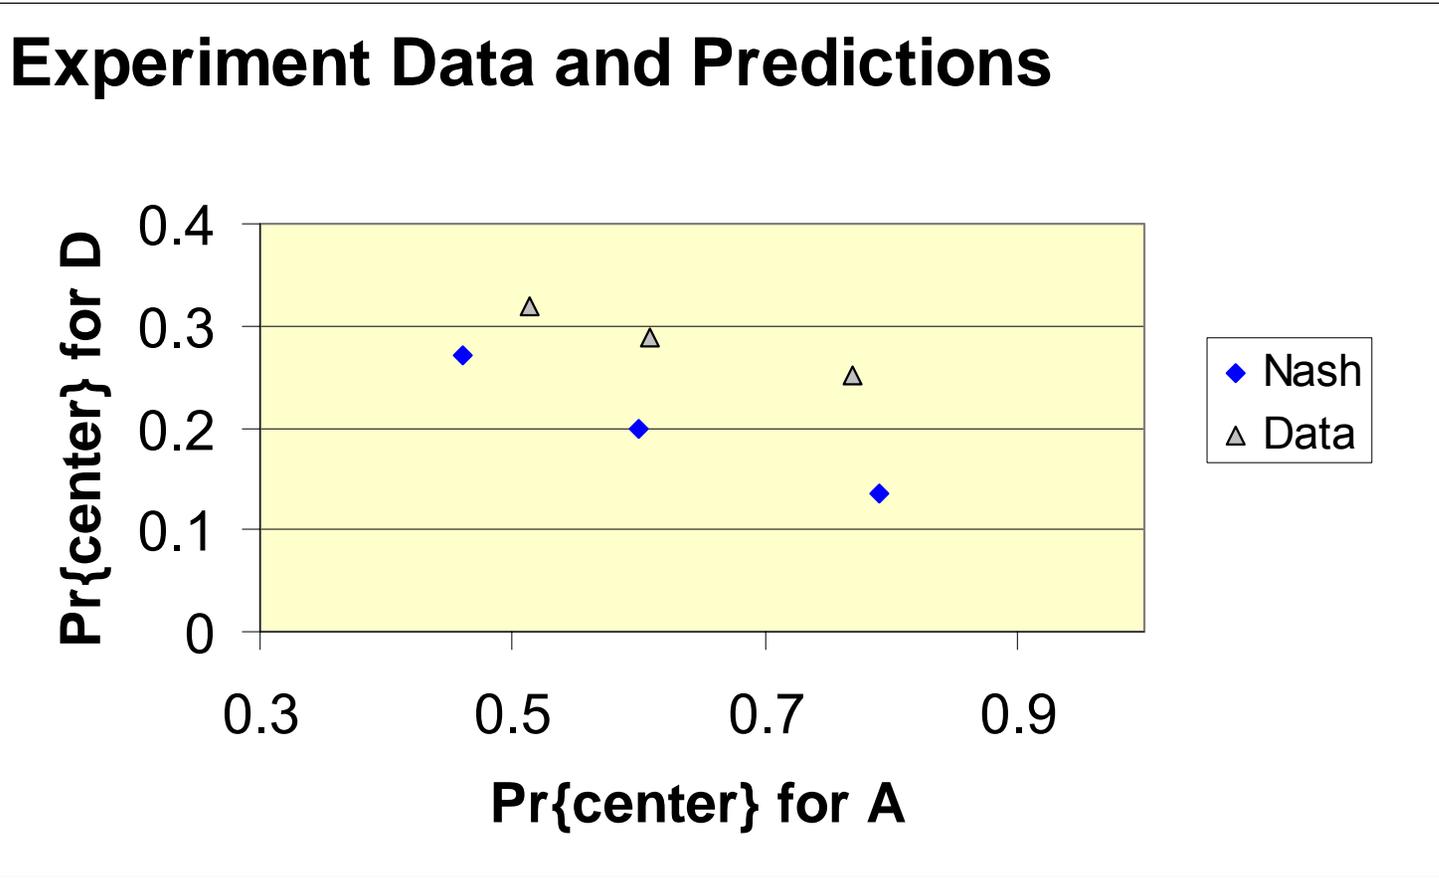


Figure 1. Experiment Data and Nash Predictions

Figure 2. Logit correspondence for Uniform treatment

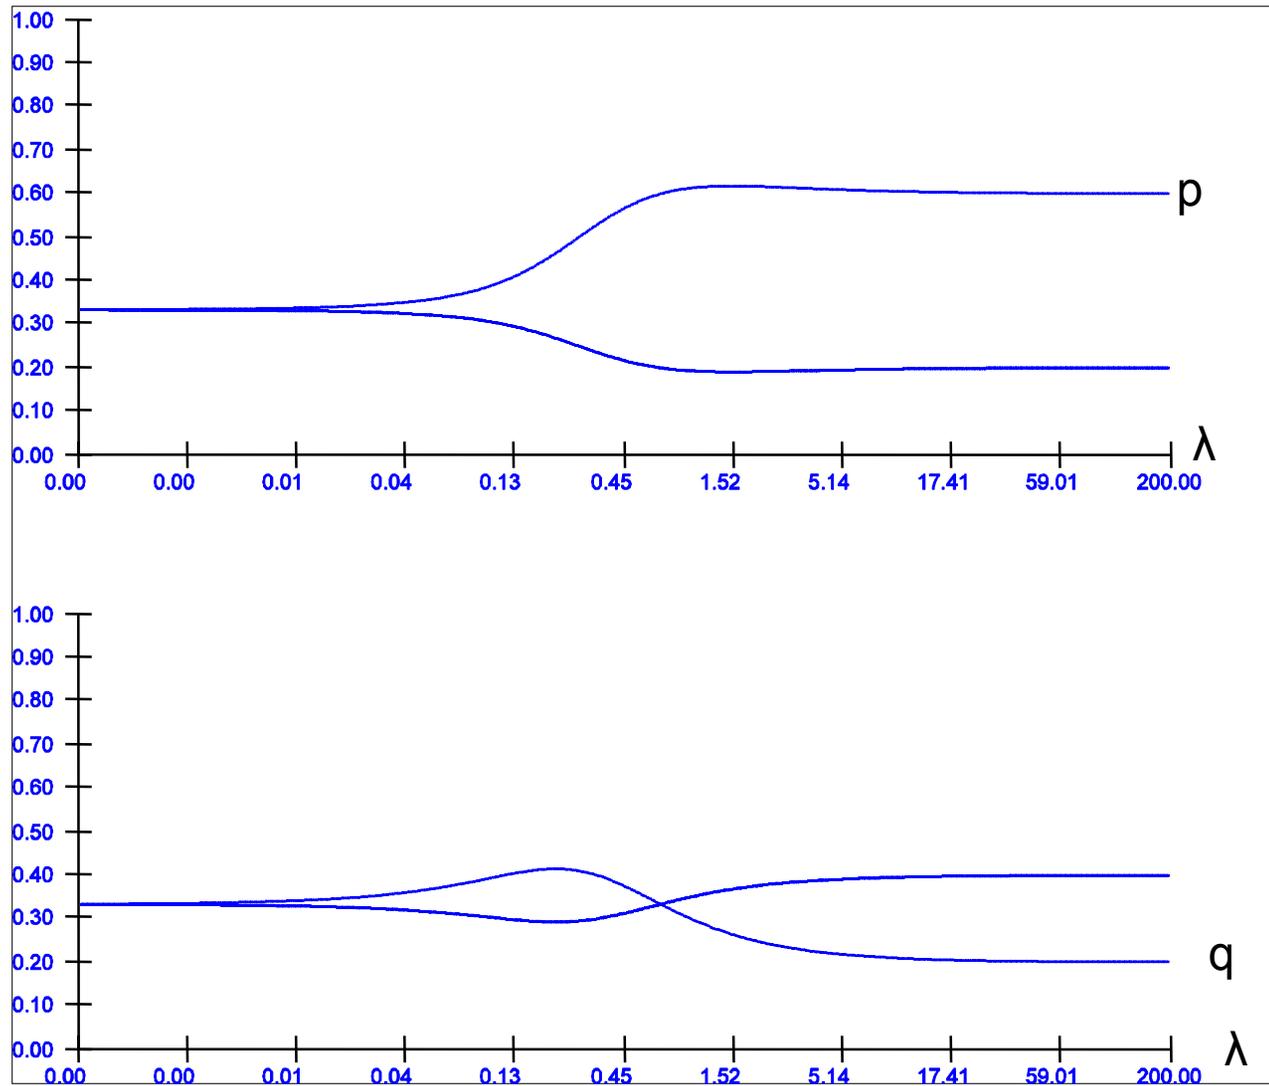


Figure 3. Logit correspondence for Low treatment

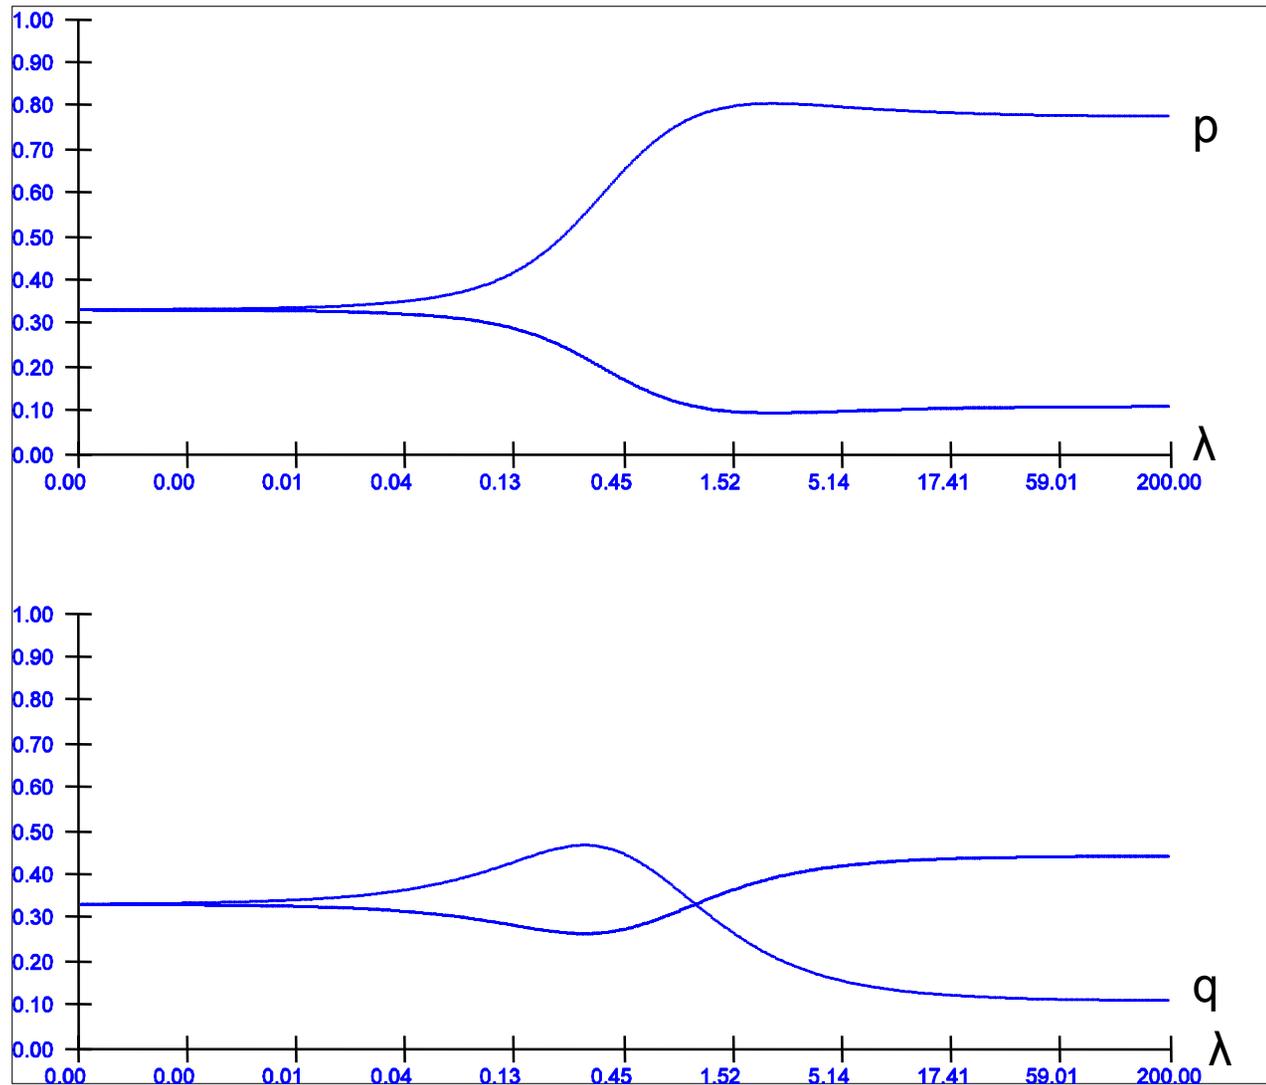
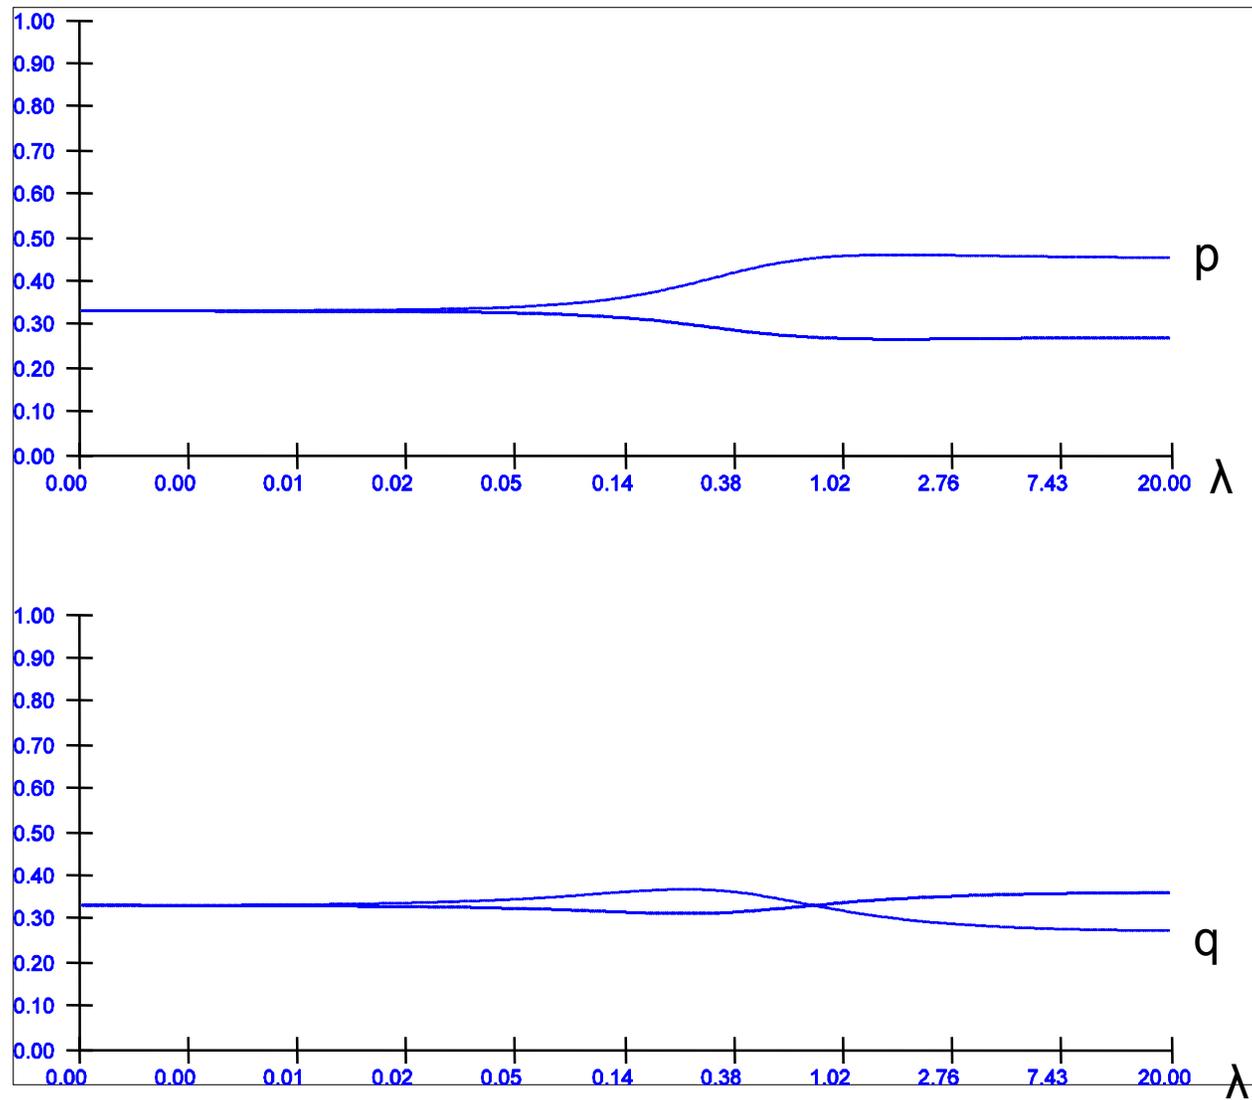


Figure 4. Logit correspondence for High treatment



Experiment Data and Predictions

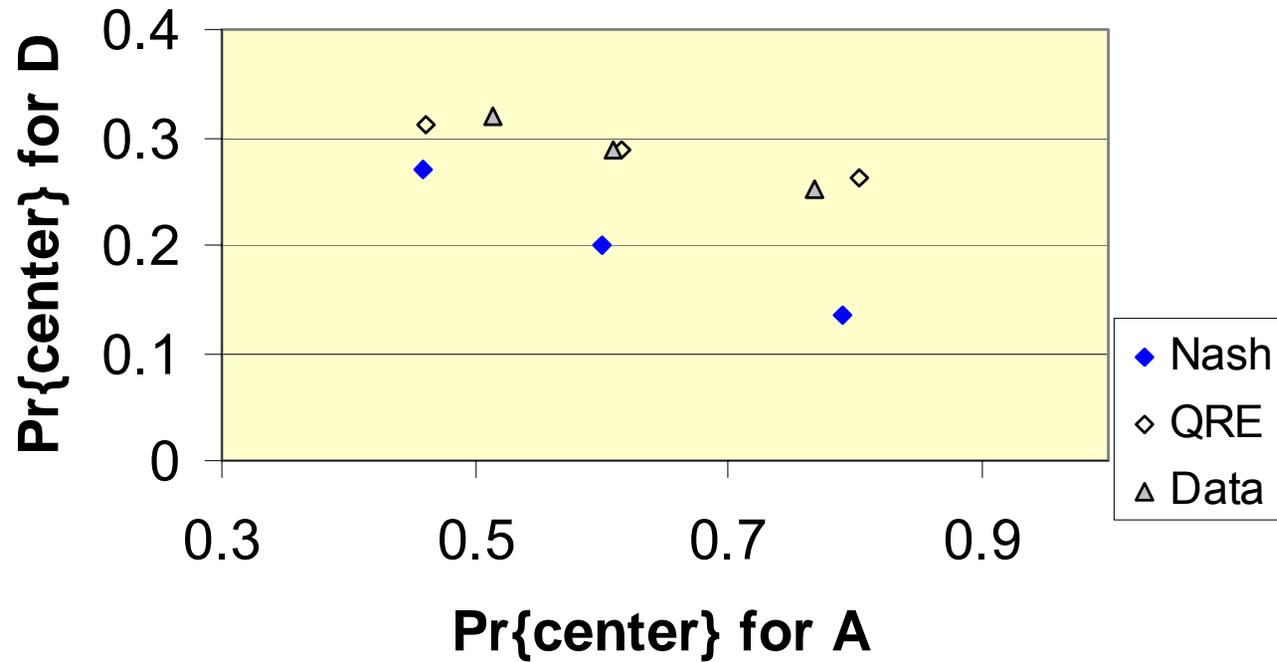


Figure 5. QRE estimates and the data.

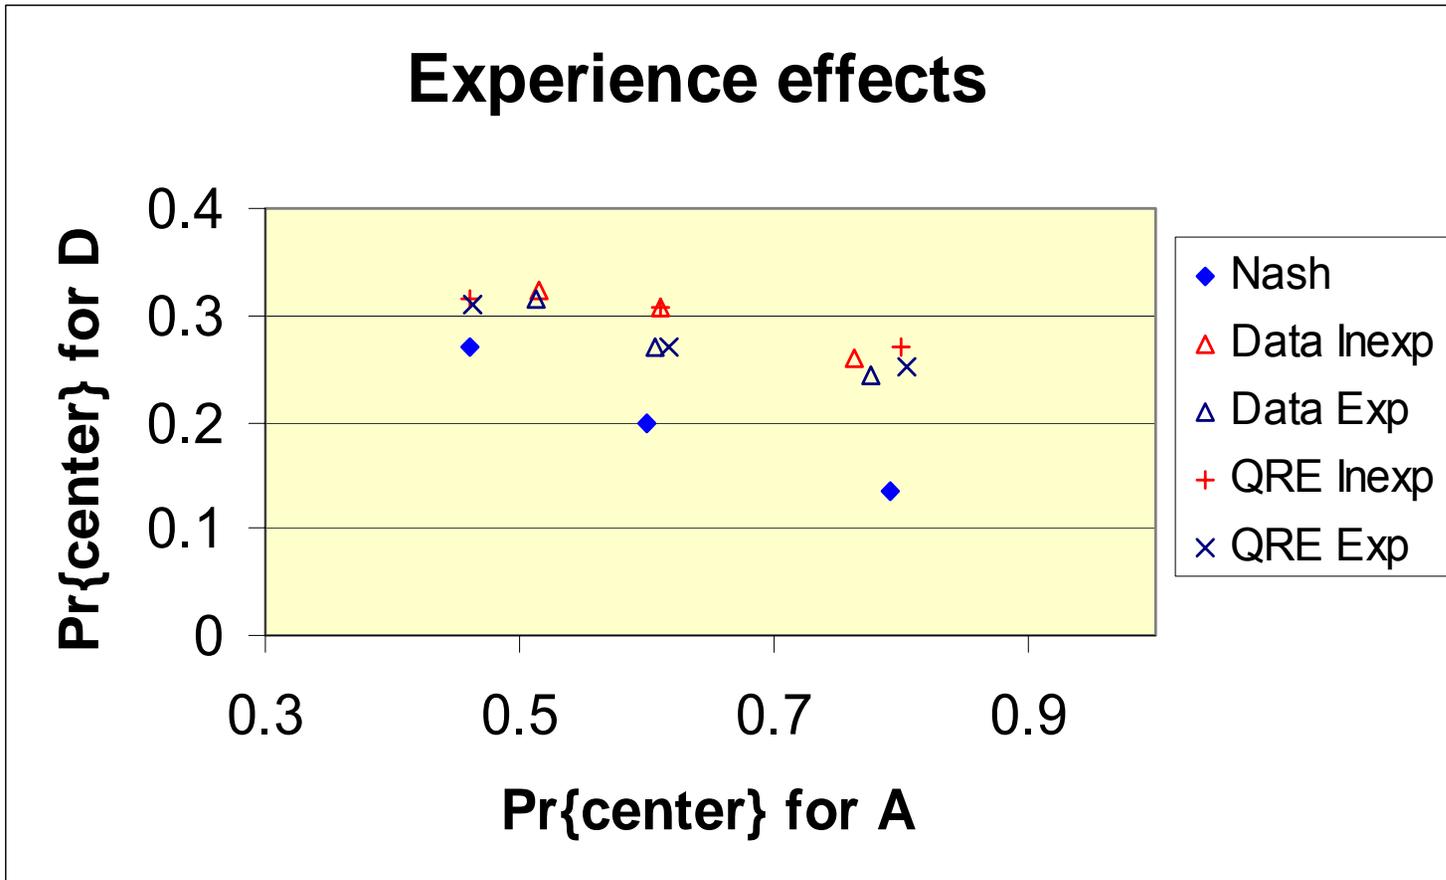


Figure 6. QRE estimates and the data. The effect of experience

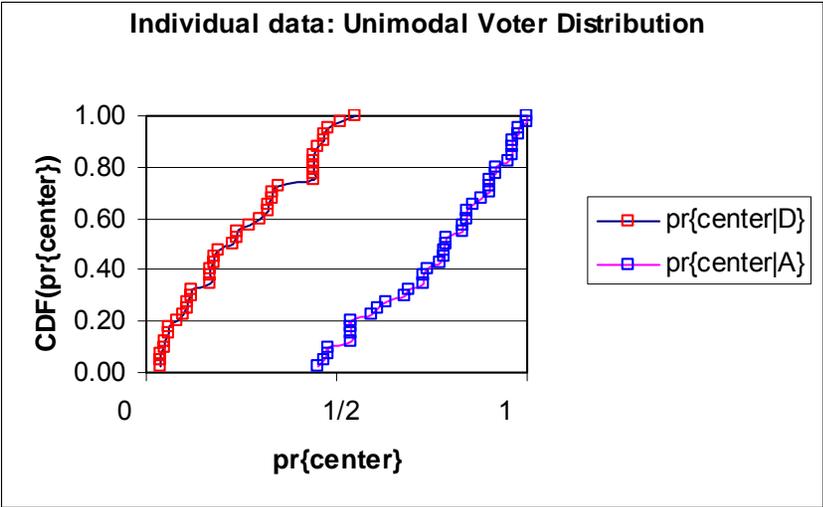
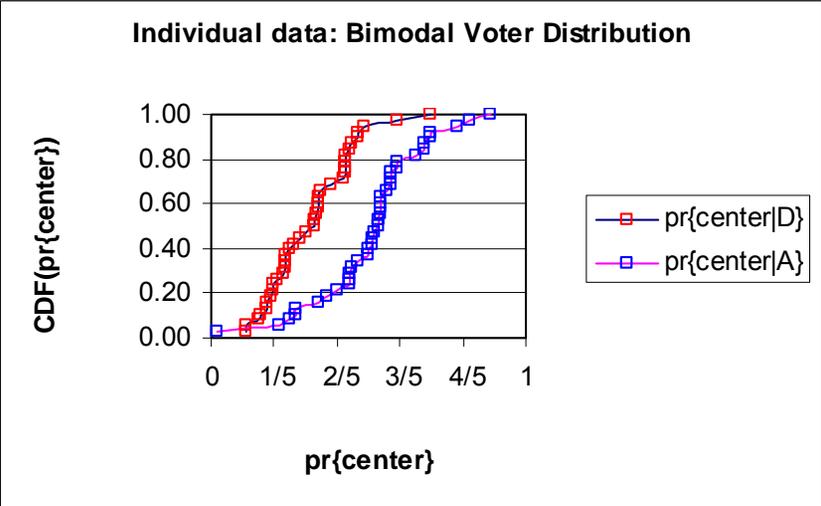
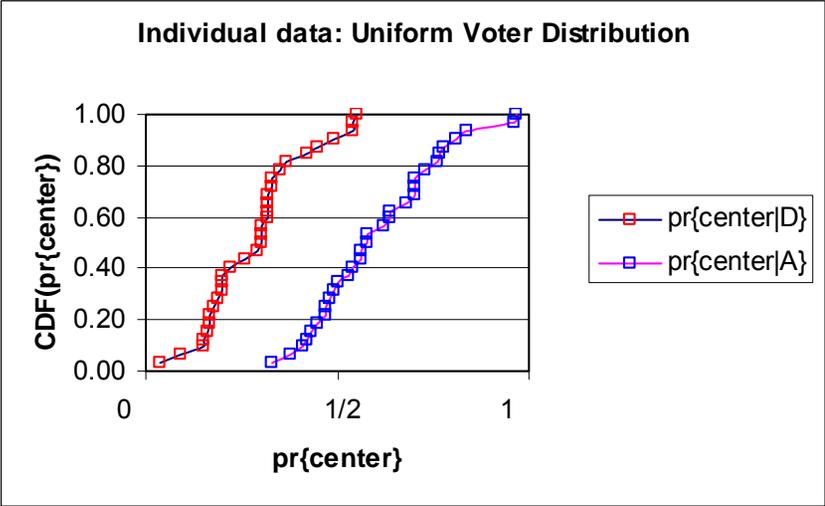


Figure 7. Choice frequencies. Individual level data.

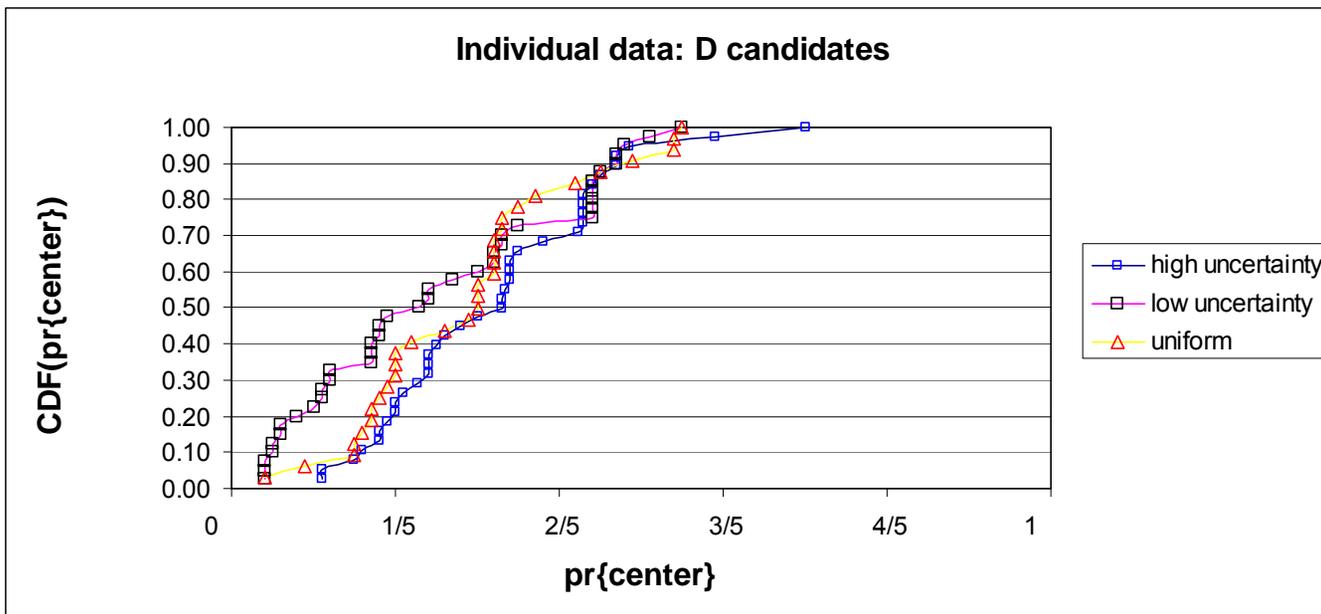
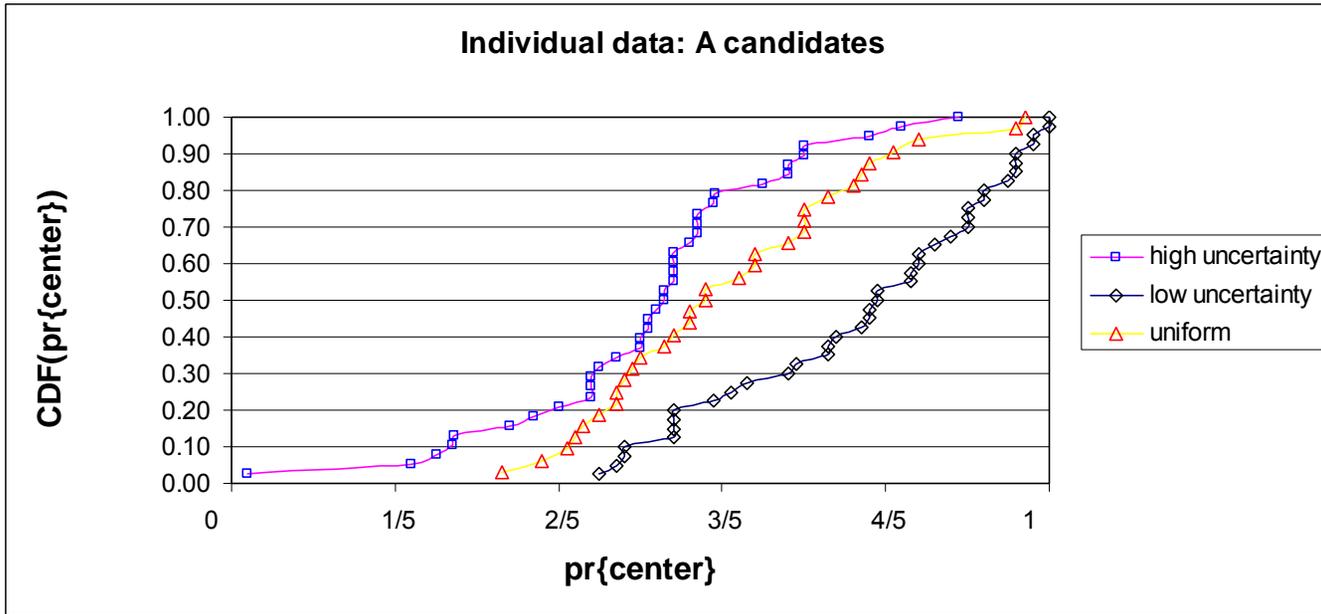


Figure 8. CDF's of Individual Choice frequencies.

Comparison of CIT and UPF data

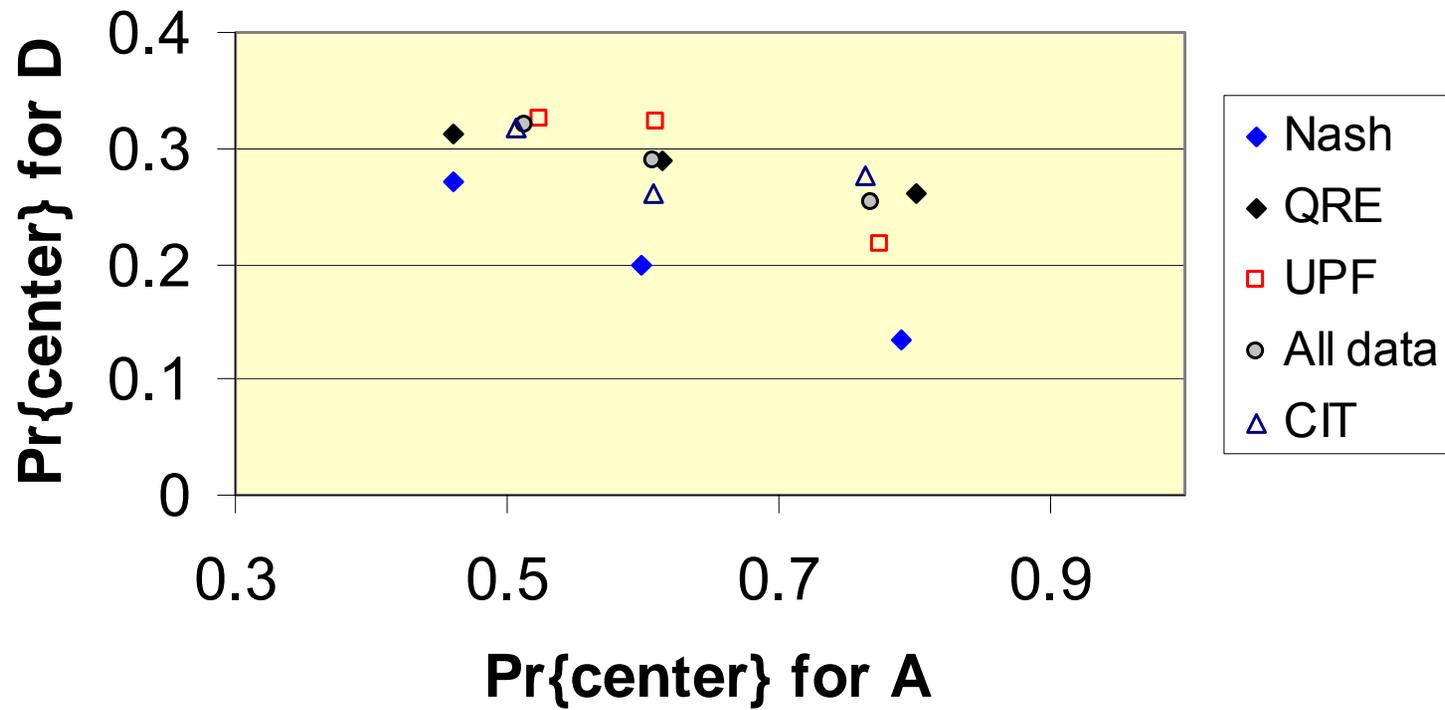


Figure 9. Choice frequencies. UPF vs. CIT.