Rational QRE: Endogenizing the Noise Parameter *

Evan Friedman†

June 24, 2016

Abstract

I modify *Quantal Response Equilibrium* (QRE) by introducing Rational Inattention. Called *Rational Quantal Response Equilibrium (Rational QRE* or *RQRE)*, the equilibrium concept microfounds, endogenizes, and *rationalizes* QRE-type noise parameters and provides a theory for the heterogeneity of noise parameters. In a Rational QRE, agents play a *dual game* where payoffs are derived from the underlying *original game*. In the first stage of the dual game, agents choose an information structure subject to information acquisition costs. In the second stage, players observe private signals from the chosen information structures about the expected utility from each of their actions given some beliefs about other players' play, and then players take actions. In equilibrium, first-stage information choice is a best response to beliefs, second-stage actions are rational given signal realizations and beliefs, and beliefs are consistent with the probability distribution of actions induced by the chosen information structures. As an empirical test, I show that Rational QRE outperforms standard QRE in the experimental data of *Palfrey, McKelvey, and Weber* (2000) even when QRE is estimated *in-sample* and Rational QRE is estimated *out-of-sample.*

*Very preliminary. I thank Alessandra Casella, Mark Dean, Han Huynh, RC Lim, Janet Lu, Lan Nguyen, Suanna Oh, Tom Palfrey, Jon Zytnick, and participants at the Columbia Economics Microtheory Colloquium for helpful comments. All errors are my own.

†Email: ekf2119@columbia.edu

1 Introduction

Quantal Response Equilibrium (QRE), first introduced by McKelvey and Palfrey [1995], has found wide application in experimental game theory. In the lab, subjects do not play Nash strategies reliably, and QRE was developed as an alternative equilibrium theory that explains what we do observe. In essence, it assumes that players make *calculation errors* in estimating expected utilities from each of their actions given beliefs of others' play, and in equilibrium, each player takes the action that gives highest noisy expected utility, and beliefs are consistent with play. Importantly, QRE introduces noise not to final Nash outcomes, but to expected utilities–hence the interpretation as calculation errors. This raises a number of questions. Where do these errors come from? Should these errors vary systematically across games? What about across players?

In a standard QRE model, the distribution of errors is completely exogenous, unrelated to the underlying game. Typically, the joint distribution of errors (for all players and all actions) is specified up to a single constant. The most common form of errors assumed in the literature is extreme value, where the parameter λ cannot be specified *a priori*. The resulting *Logit QRE* is so popular because it results in closed-form expressions for the *quantal response functions*–a fixed point in which defines the QRE. In practice the best-fit λ is estimated in-sample from data, and the best-fit QRE prediction is compared against Nash and other models.

One issue with exogenous λ is that there is no microfounded justification for how it should vary across games. Therefore, the ability of QRE to make out-of-sample predictions–estimating λ in one game and using that to make predictions about another–is severely limited. Of course, assumptions can be made that will allow QRE to make out-of-sample predictions. For instance, one might assume that λ is fixed across games. Thus, an estimate of λ from game 1 will imply very precise predictions about play in game 2, yet in practice these types of predictions are rejected.

There is also no reason to believe that λ is constant across players within a game. Heterogeneous λ_i's can be estimated–one for each player–and allowing for heterogeneity typically greatly improves fit. This of course suggests that something about the game itself implies heterogeneity, but allowing for so many degrees of freedom in estimation is problematic.

This paper attempts to address these issues by microfounding, endogenizing, and *rationalizing* the source of QRE-type calculation errors using the framework of *Rational Inattention*. The idea is that players optimally and *strategically* choose the precision of their errors subject to information acquisition costs. The equilibrium concept, called *Rational Quantal Response Equilibrim* (RQRE) is basically a Nash Equilibrium in information choice in the first stage, followed by standard QRE in the second stage using the noisiness endogenously determined. Thus, "calculation errors" is an appropriate interpretation but the errors are entirely *rational*–meaning that players may choose to make smaller errors, but it is correctly judged to be too costly to do so.

Of course, structure must be added to the form of errors and information costs, and something must be estimated from data. In this case, it is the *information costs parameter* which gives the per-unit cost of information. This parameter is arguably much more primitive than λ itself. Thus, across certain games–such as those that differ only in *payoff magnitude*–the information costs parameter is plausibly fixed, allowing the model to make strong out-of-sample predictions. In addition, Rational QRE provides a theory for the heterogeneity of λ. The theory suggests not only that λ is heterogenous, but pins down the λ_i's up to a single parameter.

This research agenda is inspired by the vast literature on QRE. In many papers, there is a discussion about where λ comes from–including the original 1995 paper. In particular, McKelvey et al. [2000] tests the QRE predictions of scaling up games by increasing the magnitudes of payoffs. That paper discusses the role of endogenizing λ, and though it stops short of a formal theory, is in much the same spirit as this work. To my knowledge, the only formal theory of endogenous λ comes from the unpublished manuscript McKelvey et al. [1997]. The approach I take is very different, but this paper can be regarded as an extension and deepening of that work.

As an empirical application, I use the experimental data from McKelvey et al. [2000] in which two games are played that differ only in payoff magnitude. I find that Rational QRE outperforms QRE when the parameters are estimated in-sample using maximum likelihood. I also find that estimating the Rational QRE parameter for one game and using it to make *out-of-sample* predictions for the other outperforms *in-sample* QRE.

2 Theory

2.1 Model

A finite, normal form game $\Gamma = \{N, A, u\}$ is defined by a set of N agents, action space $A = A_1 \times ... \times A_N$ with $A_i = \{a_{i1}, ..., a_{iJ(i)}\}$ such that each player i has $J(i)$ possible actions, and a vector of payoff functions $u = (u_1, ..., u_N)$ with $u_i : A \to \mathbb{R}$. Let ΔA_i be the set of probability measures on A_i. Elements of ΔA_i are of the form $p_i : A_i \to \mathbb{R}$ where $\sum_{a_{ij} \in A_i} p_i(a_{ij}) = 1$ and $p_i(a_{ij}) \geq 0$. Henceforth define $p_{ij} \equiv p_i(a_{ij})$. I abuse notation by using a_{ij} to refer to the action $p_i \in \Delta A_i$ with $p_{ij} = 1$. Extend payoffs to be defined over ΔA by the rule $u_i(p) = \sum_{a \in A} p(a)u(a)$. Refer to this underlying game as the *original game*.

In a Rational QRE, agents actually play a *dual game* $\Sigma = \{\Gamma, \Omega, C\}$ where Γ represents all features of the *original game*, $\Omega = \Omega_1 \times ... \times \Omega_N$ is a space of information structures with $\omega_i \in \Omega_i$ a typical element, and $C \equiv (c_1, ..., c_N)$ are *information cost* functions with $c_i : \Omega_i \to \mathbb{R}$. Define $\Delta\Omega_i$ to be the set of probability measures on Ω_i with $\sigma_i \in \Delta\Omega_i$ a typical element. Extend c_i to be defined on $\Delta\Omega_i$ by $c_i(\sigma_i) = \sum_{\omega_i \in \Omega_i} \sigma_i(\omega_i)c_i(\omega_i)$. The dual game has *two stages* with timing as follows:

1. *Information acquisition.* Each player i chooses an information structure $\sigma_i \in \Delta\Omega_i$ and pays costs $c_i(\sigma_i)$.

2. *Playing actions.* Each player i observes private signals from σ_i and takes an action $p_i \in \Delta A_i$.

2.2 Notation and Assumptions

Beliefs. Denote by $\beta \in \Delta A$ the profile of beliefs of all players' actions, $\beta_{-i} \in \Delta A_{-i}$ the beliefs of the actions of all players other than i, and β_{ij} the belief that i will play action j.

Utility. Define $u_{ij}(\beta_{-i}) \equiv \mathbb{E}_{a_{-i}}[u_i(a_{ij}, a_{-i})|\beta_{-i}]$ to be the "true" expected utility of player i from taking action a_{ij} *given beliefs* of other players' actions. Only if chosen information structures are perfectly informative will players observe this object without error (which will never occur in equilibrium, as will become clear).

Information Structures. Assume that all information structures can be equivalently expressed in terms of independent and identically distributed *noise* around the true expected utilities (given beliefs). Assume symmetric information spaces $\Omega_i \equiv \Omega = [0, \infty)$ for all i where $\omega_i \in \Omega$ corresponds to some notion of the "precision" of $J(i)$ independent draws $\epsilon_i \equiv (\epsilon_{i1}, ..., \epsilon_{iJ(i)})$. I will further assume one of two mutually exclusive sets of restrictions on the errors. Unless otherwise stated, the errors are assumed to satisfy *exactly one* of the following:

(A1). ϵ_{ij} is distributed i.i.d. according to a normal distribution with mean zero and precision ω_i. That is, $\epsilon_{ij} \sim_{iid} N(0, \frac{1}{\omega_i})$. Call the resulting Rational QRE a "Probit Rational QRE."

(A2). ϵ_{ij} is distributed i.i.d. according to a mean-zero extreme value (Gumbel) distribution with CDF $F(\epsilon_{ij}) = e^{-e^{-\omega_i \epsilon_{ij} - \gamma}}$ where $\gamma \approx 0.5772$ is the Euler-Mascheroni constant. These are the familiar *logit errors.* Call the resulting Rational QRE a "Logit Rational QRE."

Information Costs. Assume that the information costs function is given by $c_i(\omega_i) = \rho_i k(\omega_i)$, where $\rho_i > 0$ is i's *information cost parameter* and $k^{'}(\cdot) > 0$, $k^{''}(\cdot) > 0$, $k^{'}(0) = 0$, and $lim_{\omega_i \to \infty} k^{'}(\omega_i) = \infty$. In particular, quadratic costs $c_i(\omega_i) = \rho_i \omega_i^2$ satisfy these restrictions. Note that ρ_i is the parameter that must be estimated from data to fit the model.

2.3 Equilibrium Concept

A Rational QRE consists of *information choice strategies* $\{\sigma_i^*(\beta_{-i}; \rho_i)\}_i$ and *action strategies* $\{p_i^*(\beta_{-i}, \epsilon_i)\}_i$ such that

(S1). *First-stage information choice best response.* For all i, $\sigma_i^*(\beta_{-i}; \rho_i) \in \Delta\Omega_i$ is such that all ω_i with $\sigma_i^*(\omega_i) > 0$ solve

$$\underset{\hat{\omega}_i \in [0,\infty)}{max} \ \mathbb{E}_a[u_i(a)|\hat{\omega}_i, \beta_{-i}] - \rho_i k(\hat{\omega}_i)$$

(S2). *Second-stage rationality.* For all i, $p_i^*(\beta_{-i}, \epsilon_i) \in \Delta A_i$ is such that all j with $p_i^*(a_{ij}) > 0$ satisfy

$$u_{ij}(\beta_{-i}) + \epsilon_{ij} \geq u_{ik}(\beta_{-i}) + \epsilon_{ik} \text{ for all } k$$

(S3). *Consistency of beliefs.* $\beta \in \Delta A$ is such that the ex-ante probability of realized actions coincides with β:

$$\mathbb{P}(a_i^* = a_{ij}) = \beta_{ij} \text{ for all } i \text{ and } j$$

Pure versus mixed. A Rational QRE is *pure* if, for all players, σ_i^* is degenerate, i.e. all players choose a single information structure with probability one, and *mixed* otherwise. Since actions are chosen conditional on signal realizations, players can only mix over actions if signal realizations are such that they are indifferent to several actions. However, the restrictions introduced in this section imply that this cannot happen with positive probability in equilibrium, so the characterization of pure and mixed strategy equilibria in terms of *information choice only* is without loss.

2.4 Scale

Standard QRE gives very precise predictions about what happens when games are "scaled up" when the noise parameters are assumed fixed. In this case, scaling up a game by multiplying all payoffs by a positive constant greater than 1 is equivalent to reducing the magnitude of errors. At first glance, this seems a reasonable implication. However, under the classic interpretation of errors as stemming from calculation mistakes, why should this be so? Is it more difficult to calculate the average of 5 and 10 cents than it is to calculate the average of 5 and 10 dollars? Surely, under a fixed attention model, one would expect the errors from both calculations to have the same distribution when expressed in terms of cents and dollars respectively.

Of course, if attention is flexible and costly, and the payoffs are related to scale, we would expect a very different outcome. Consider the following decision problems in which a *risk-neutral* decision maker must choose a lottery (A or B) from menues (1 or 2):

$$\text{Menu 1:} \quad \overset{A}{\left\{ \begin{array}{c} \$0.05 \text{ with prob. } 0.5 \\ \$0.10 \text{ with prob. } 0.5 \end{array} \right\}} \quad \overset{B}{\{\$0.07 \text{ with prob. } 1\}}$$

$$\text{Menu 2:} \quad \overset{A}{\left\{ \begin{array}{c} \$5 \text{ with prob. } 0.5 \\ \$10 \text{ with prob. } 0.5 \end{array} \right\}} \quad \overset{B}{\{\$7 \text{ with prob. } 1\}}$$

The lotteries in menues 1 and 2 differ *only in payoffs*, so the complexity of choice is held fixed. Under a fixed attention model, the frequency of choice mistakes should be the same under both menues. However, if attention is flexible and costly, we would expect that at optimum, the decision maker would pay more attention under menu 2 than under menu 1 and hence choose option A–the higher expected value lottery–more often.

Examples like these motivate an equilibrium model in which scale does effect the prevalence of choice mistakes, but only through endogenous information choice and equilibrium effects–not through the equivalent of exogenously varying the magnitude of errors as in standard QRE. Rational QRE offers a natural way of incorporating this idea. Simply define $\chi > 0$ to be the "scale of the game" and replace requirement **(S1)** with

(S1$'$). *First-stage information choice best response.* For all i, $\sigma_i^*(\beta_{-i};\rho_i,\chi) \in \Delta\Omega_i$ is such that all ω_i with $\sigma_i^*(\omega_i) > 0$ solve

$$\max_{\hat{\omega}_i \in [0,\infty)} \chi \mathbb{E}_a[u_i(a)|\hat{\omega}_i,\beta_{-i}] - \rho_i k(\hat{\omega}_i)$$

To analyze a single game in isolation, set $\chi = 1$ without loss, but having χ appear explicitly in the first stage (and the *first-stage only*) allows for interesting *across game* comparisons not possible with standard QRE. To see this, first consider the following definition and assumption:

Definition 1. Given a normal form game $\Gamma = \{N, A, u\}$, define the **scale family** of Γ, denoted $\mathcal{F}(\Gamma)$, as follows: $\Gamma' = \{N', A', u'\} \in \mathcal{F}(\Gamma)$ if and only if $N = N' \equiv \bar{N}$, $A = A' \equiv \bar{A}$ and there exists $\chi > 0$ such that $u' = \chi u$ (i.e. for all $a \in \bar{A}$ and $i \in \bar{N}$, $u_i'(a) = \chi u_i(a)$) in which case identify Γ' with χ.

Assumption 1. *Fix $\Gamma = \{N, A, u\}$. The Rational QRE of all games $\chi \in \mathcal{F}(\Gamma)$ are defined using the same payoff function u and game specific-scale χ, i.e. the only difference in Rational QRE between games in $\mathcal{F}(\Gamma)$ comes from the χ that defines them.*

With this assumption in hand, we get automatically that all games in $\mathcal{F}(\Gamma)$ have the fundamental property that relative scale will effect play *only* through the information choice (and equilibrium effects). That is, *conditional on information choice*, χ has absolutely *no effect* on play, but χ does effect the incentives to acquire information (as well as the formation of beliefs in equilibrium). This is equivalent to the structural observation that χ scales payoffs as well as *errors*, and hence χ "drops out" of the second stage. Note that this is a major departure from standard QRE model, in which scaling both errors and payoffs has no effect on equilibrium outcomes.

Note also that χ is isomorphic to ρ: for any χ and ρ, one can eliminate χ by dividing the first-stage objective by χ and defining $\tilde{\rho} \equiv \frac{\rho}{\chi}$ without changing the resulting equilibrium. Therefore, when estimating ρ *in-sample* via maximum likelihood, the value of χ has no effect on model fit or predictions, but it does effect the associated estimate $\hat{\rho}$. Therefore, χ can be set arbitrarily to analyze a single game isolation, but setting *relative* χ correctly is essential for *out-of-sample* predictions–estimating $\hat{\rho}$ for one game and using that to make predictions about a scaled version of the same game. Keeping χ as an explicit parameter is also conceptually appealing, because unlike ρ, which is assumed to come exogenously from within the players (from the brain perhaps), χ can be varied experimentally and is known to researchers in the lab.

2.5 Interpretation

Big Picture. When presented with a game, players must exert effort to understand any aspects of the game. This includes the basic rules of the game and the payoffs to their actions. It is cognitively costly to understand the game completely and juggle in the mind the payoffs to each action.

These "cognitive costs" are intuitive as anecdotally people avoid thinking harder than they have to, and they are empirically justified by experimental decision problems in the lab (see for example Caplin and Dean [2015]). Besides the *costs* of "attention," there is the obvious *benefit*. When paying close attention, there will be fewer and less costly calculation errors. Balancing these tradeoffs implies an optimal or rational level of attention–or equivalently *inattention*.

Because of these tradeoffs, it is intuitive that players will respond to the stakes of the game. In decision problems, it is obvious that increasing the stakes will lead to greater attention at optimum. In games, however, it is more ambiguous. Comparing two games of differing stakes, player i will be effected in *two ways* in going from a low stakes game to a high stakes game. First, there is the direct effect. When the stakes are higher, mistakes are more costly, so player i will be willing to pay more attention for any given beliefs of others' play (like in a decision problem). Second, there is the equilibrium effect. Player i knows that other players will similarly respond to the stakes, so beliefs too will change in the new equilibrium.

Interpretation Issues. The Rational QRE concept implies players have a very particular understanding of a game, with perfect information about some aspects of the game and imperfect information about others.

Rational QRE posits that players have a completely objective (as opposed to subjective) payoff function that they observe without error and wish to maximize in expectation, i.e. the "true" utility. Players are also forward looking in the sense that they know that they will make mistakes when it comes to taking actions. Thus, in stage 1, they make a *costly investment* in their level of attention under perfect information, and in stage 2, they still make choice mistakes as if some information has been lost.

From an economic perspective, these assumptions seem at odds. The information required to make the optimal investment in attention is exactly the same as is required to take the optimal action, so how is this a sensible model?

One might argue that the investment in attention is a fundamentally different problem than action choice. If that were the case, then attention in investment would be subject to errors that cannot be derived from the same errors that contribute to choice mistakes. If one accepts this, the model is justified as an approximation to the case where attention investment errors are negligible. Of course, one could write down a model where there is exogenous noise preventing optimal attention investment, but that would add (at least) one additional parameter that needs to be estimated.

Another interpretation that I do not take here, is simply that the choice mistakes are not a function of information at all, and are just "trembles." Under this assumption, players always have perfect information, but are just sloppy in executing strategies. They know what the best action is, but they just happen to press the wrong button or misspeak. They can pay more attention to reduce these types of mistakes, but no learning takes place.

Rational Inattention. Whether or not the information choice constitutes "rational inattention" is a subject of debate. The primary reason is that there has been no discussion of priors and states. To shoehorn the model into a Rational Inattention framework, consider as states the expected utilities for all actions given beliefs. The specification with normal errors results from improper uniform priors on the real line and normal signals centered around true expected utilities.

2.6 Existence

Theorem 1. *In the general $N \times 2$ game, there exists a* pure strategy *Rational QRE.*

Proof. In the *second stage*, given beliefs β_{-i} and the information structure chosen in the first stage ω_i, i will play a_{i1} if and only if

$$\begin{aligned} u_{i1}(\beta_{-i}) + \epsilon_{i1} &\geq u_{i2}(\beta_{-i}) + \epsilon_{i2} \iff \\ \Delta_i(\beta_{-i}) &\geq \tilde{\epsilon}_i \end{aligned}$$

where $\Delta_i(\beta_{-i}) \equiv u_{i1}(\beta_{-i}) - u_{i2}(\beta_{-i})$ is the difference in expected utility between actions, and $\tilde{\epsilon}_i \equiv [\epsilon_{i2} - \epsilon_{i1}]$ is a composite error term. Under **(A1)**, $\tilde{\epsilon}_i \sim N(0, \frac{2}{\omega_i})$ and the probability of playing a_{1i} and a_{2i} are given by

$$\begin{aligned} Q^{(1)}_{i1}(\omega_i, \beta_{-i}) &\equiv \mathbb{P}(\tilde{\epsilon}_i \leq \Delta_i(\beta_{-i})) \\ &= \Phi\left(\sqrt{\frac{\omega_i}{2}}\Delta_i(\beta_{-i})\right) \\ Q^{(1)}_{i2}(\omega_i, \beta_{-i}) &= 1 - \Phi\left(\sqrt{\frac{\omega_i}{2}}\Delta_i(\beta_{-i})\right) \end{aligned}$$

where Φ is the standard normal CDF. Under **(A2)**, the analagous expressions are given by the standard *logit quantal response function* (see for example McFadden [1976]):

$$\begin{aligned} Q^{(2)}_{ij}(\omega_i, \beta_{-i}) &= \frac{e^{\omega_i u_{ij}(\beta_{-i})}}{\sum_{k=1}^{2} e^{\omega_i u_{ik}(\beta_{-i})}} \\ &= \frac{1}{1 + e^{-\omega_i \Delta_i(\beta_{-i})}} \text{ for } j \in \{1, 2\} \end{aligned}$$

In the *first stage*, i chooses ω_i, the precision of an information structure, to maximize expected utility net of costs, i.e. for each type of error $g \in \{1, 2\}$, i solves

$$\underset{\omega_i}{max}\ \underbrace{\chi \sum_{k=1}^{2} Q_{ik}^{(g)}(\omega_i, \beta_{-i}) u_{ik}(\beta_{-i})}_{\substack{\text{information}\\ \text{benefit}}} - \underbrace{\rho_i k(\omega_i)}_{\substack{\text{information}\\ \text{cost}}}$$

Define the optimal information choice as the solution to this problem

$$\omega_i^{(g)*}(\beta_{-i}; \chi, \rho_i) \equiv \underset{\omega_i}{argmax}\ \chi \sum_{k=1}^{2} Q_{ik}^{(g)}(\omega_i, \beta_{-i}) u_{ik}(\beta_{-i}) - \rho_i k(\omega_i) \tag{1}$$

By Lemma 1 (see Appendix 6.1), $\omega_i^{(g)*}(\beta_{-i}; \chi, \rho_i)$ exists, is unique, and is continuous in $(\beta_{-i}, \chi, \rho_i)$. Thus, the probability of playing a_{i1} given beliefs is

$$\mathbb{P}_{i1}^{(g)}(\beta_{-i}; \chi, \rho) \equiv Q_{i1}^{(g)}(\omega_i^{(g)*}(\beta_{-i}; \chi, \rho_i), \beta_{-i}) \tag{2}$$

(1) and (2) give both *first-stage information choice best response* and *second-stage rationality* given beliefs. All that remains for a Rational QRE is *consistency of beliefs*, which holds if in addition we have

$$\beta_{i1} = \mathbb{P}_{i1}^{(g)}(\beta_{-i}; \chi, \rho_i) \tag{3}$$

for all i. Combining (1),(2), and (3) for each i and stacking defines a fixed point equation whose solution defines an equilibrium:

$$\beta \equiv \begin{bmatrix} \vdots \\ \beta_{i1} \\ \vdots \end{bmatrix} = \begin{bmatrix} \vdots \\ Q_{i1}^{(g)}(\omega_i^{(g)*}(\beta_{-i}; \chi, \rho_i), \beta_{-i}) \\ \vdots \end{bmatrix} \equiv f(\beta) \tag{4}$$

where β is now redefined to be the N-vector of probabilities of i playing a_{i1}. $f : [0, 1]^N \mapsto [0, 1]^N$ is continuous in β as $Q_{i1}^{(g)}$ is continuous and composed with $\omega_i^{(g)*}$ that is also continuous, again by Lemma 1. By compactness of $[0, 1]^N$, the conditions of Brouwer's fixed point theorem are satisfied which guarantees the existence of a solution. □

More than 2 actions. Unfortunately, existence is difficult to prove, even in mixed strategies, for games in which players have more than 2 actions. The "information benefit" easily generalizes to the case of an arbitrary number of actions:

$$\chi \sum_{k=1}^{J(i)} Q_{ik}^{(g)}(\omega_i, \beta_{-i}) u_{ik}(\beta_{-i})$$

However, Lemma 1 showing continuity of $\omega_i^{(g)*}(\beta_{-i}; \chi, \rho_i)$ in β_{-i} breaks down when a player has more than 2 actions. Lemma 1 rests on showing that information benefit is *concave* in ω_i for any β_{-i}. This fails in the more general case, so we cannot say that $\omega_i^{(g)*}(\beta_{-i}; \chi, \rho_i)$ is continuous, which is required for the fixed point.

2×2 games. Rational QRE cannot generally be found in closed form, but it is straightforward to solve for equilibria numerically, especially in the case of 2×2 games, which may be useful in applications. The following corollary specializes the existence result to 2×2 games.

Corollary 1. *(i) In the* 2×2 *game, a* Probit Rational QRE *is defined by*

$$\begin{aligned} \beta_{i1} &= \Phi\left(\sqrt{\frac{\omega_i^{(1)*}\left(\Phi\left(\sqrt{\frac{\omega_{-i}^{(1)*}(\beta_{i1})}{2}}\Delta_{-i}(\beta_{i1})\right)\right)}{2}}\Delta_i\left(\Phi\left(\sqrt{\frac{\omega_{-i}^{(1)*}(\beta_{i1})}{2}}\Delta_{-i}(\beta_{i1})\right)\right)\right) \\ \beta_{i2} &= 1 - \beta_{i1} \end{aligned} \tag{5}$$

for $i, -i \in \{1, 2\}$, *where*

$$\Delta_i(\beta_{-i1}) \equiv u_{i1}(\beta_{-i1}) - u_{i2}(\beta_{-i1})$$

and

$$\omega_i^{(1)*}(\beta_{-i1}; \chi, \rho) \equiv \underset{\omega_i}{argmax}\ \chi\Phi\left(\sqrt{\frac{\omega_i}{2}}\Delta_i(\beta_{-i1})\right)\Delta_i(\beta_{-i1}) - \rho_i k(\omega_i)$$

(ii) In the 2×2 *game, a* Logit Rational QRE *is defined by*

$$\begin{aligned} \beta_{i1} &= \frac{1}{1 + e^{-\omega_i^{(2)*}\left(\frac{1}{1+e^{-\omega_{-i}^{(2)*}(\beta_{i1})\Delta_{-i}(\beta_{i1})}}\right)\Delta_i\left(\frac{1}{1+e^{-\omega_{-i}^{(2)*}(\beta_{i1})\Delta_{-i}(\beta_{i1})}}\right)}} \\ \beta_{i2} &= 1 - \beta_{i1} \end{aligned} \tag{6}$$

for $i, -i \in \{1, 2\}$, *where*

$$\omega_i^{(2)*}(\beta_{-i};\chi,\rho) \equiv \underset{\omega_i}{argmax}\ \chi\frac{u_{i1}(\beta_{-i})e^{\omega_i u_{i1}(\beta_{-i})}+u_{i2}(\beta_{-i})e^{\omega_i u_{i2}(\beta_{-i})}}{\sum_{k=1}^{2}e^{\omega_i u_{ik}(\beta_{-i})}} - \rho_i k(\omega_i)$$
$$= \underset{\omega_i}{argmax}\ \chi\frac{u_{i1}(\beta_{-i})+u_{i2}(\beta_{-i})e^{-\omega_i\Delta_i(\beta_{-i})}}{1+e^{-\omega_i\Delta_i(\beta_{-i})}} - \rho_i k(\omega_i)$$

Proof. These are just special cases of (4) (found by setting $N = 2$), which by Theorem 1, characterizes the equilibria. □

2.7 Scale Effects

The results of this section make use of Defintion 1 and Assumption 1, and give testable implications of varying the scale parameter χ. Importantly, χ can be varied experimentally, which is the focus of Section 3.

The first result explains an important relationship between Rational QRE and Nash and is analagous to a result for standard QRE.

Theorem 2. *Fix a normal form game* Γ *and information cost parameters* $\{\rho_i\}_i$*. Let* $\{\chi^1,\chi^2,...\}$ *be a sequence of scale parameters such that* $lim_{t\to\infty}\chi^t = \infty$ *and let* $\{\beta^1,\beta^2,...\}$ *and* $\{\omega_i^1,\omega_i^2,...\}_i$ *be corresponding sequences of pure strategy Rational QRE mixing probabilities and information choices of* $\chi^t \in \mathcal{F}(\Gamma)$ *such that* $lim_{t\to\infty}\beta^t = \beta^*$*. Then* β^* *is a Nash equilibrium of* Γ*.*

Proof. Suppose β^* is not a Nash equilibrium. Then there is some player i and a pair of actions a_{ij} and a_{ik} such that $\beta^*(a_{ik}) > 0$ and $u_i(a_{ij},\beta^*_{-i}) > u_i(a_{ik},\beta^*_{-i})$ or equivalently $u_{ij}(\beta^*_{-i}) > u_{ik}(\beta^*_{-i})$. Since u is continuous, it follows that for sufficiently small $\epsilon > 0$ there is a T such that $u_{ij}(\beta^t_{-i}) > u_{ik}(\beta^t_{-i})+\epsilon$ for all $t \geq T$. But then since $\chi^t \to \infty$, choice mistakes become infinitely costly and any finite $\bar{\omega}_i$ is such that there exists $T^{'} \geq T$ such that $\bar{\omega}_i$ cannot be a first-stage information choice best response for any $t \geq T^{'}$. Thus, $\omega_i^t \to \infty$ and $\beta^t(a_{ik}) \to 0$, contradicting that $\beta^*(a_{ik}) > 0$. □

The next conjecture is one of the main comparative statics results and is easily falsifiable in data. It basically says that when games are scaled up, you can always find equilibria characterized by information structures with greater precision.

Conjecture 1. *Fix an* $N\times 2$ *normal form game* Γ*. Define* $\{\omega_i^{e*}(\chi)\}_{ie}$ *to be the set of all pure strategy Rational QRE information choices for game* $\chi \in \mathcal{F}(\Gamma)$ *in equilibrium e (allowing for multiplicity). Take any* $\chi,\chi^{'} \in \mathcal{F}(\Gamma)$ *with* $\chi < \chi^{'}$*. There exists a Rational QRE of* $\chi^{'}$*,* $\bar{e}$*, in which for all e and i,* $\omega_i^{\bar{e}*}(\chi^{'}) \geq \omega_i^{e*}(\chi)$*.*

Proof. ... □

3 Empirical Application

For this section, further specialize the notation for 2×2 games. Let $U \equiv a_{11}$, $D \equiv a_{12}$, $L \equiv a_{21}$, and $R \equiv a_{22}$ for "Up," "Down," "Right," and "Left" so that $A_1 = \{U, D\}$ and $A_2 = \{L, R\}$. Also define $\alpha \equiv \beta_{11}$ and $\beta \equiv \beta_{21}$ to be the probabilities that player 1 plays U and player 2 plays L respectively.

As an empirical application, I revisit the experiment of McKelvey et al. [2000] which explores the effect of payoff magnitude. Consider their "Asymmetric Matching Pennies" games, in particular what they call games A and C as described in Table 1. Importantly, note that C is a scaled up version of A, i.e. $A, C \in \mathcal{F}(A)$ with A and C corresponding to $\chi = 1$ and $\chi = 4$ respectively.

player 2

A	L	R
U	$9, 0$	$0, 1$
D	$0, 1$	$1, 0$

C	L	R
U	$36, 0$	$0, 4$
D	$0, 4$	$4, 0$

player 1

	α	β	n
A	0.643	0.241	1800
C	0.594	0.257	1200

Table 1: *Asymmetric Matching Pennies from* McKelvey et al. [2000]. This table gives payoff matrices and play data of games A and C. In their experiment, each unit of payoff corresponded to \$0.10.

Estimation. I estimate *Probit Rational QRE* (*Logit* gives similar estimates) with symmetric, quadratic information costs $c_i(\omega) = c_j(\omega) \equiv c(\omega) = \rho\omega^2$ for all players i and j. Rational QRE is calculated numerically using the fixed point equation (5), and $\hat{\rho}$ is chosen to maximize the likelihood of the data. The results of fitting standard Logit QRE and Rational QRE to games A and C via maximum likelihood are presented in Table 2.

Game	QRE	QRE	QRE	**RQRE**	**RQRE**	**RQRE**	**RQRE**	**RQRE**	QRE	**RQRE**
	$\hat{\alpha}$	$\hat{\beta}$	$\hat{\lambda}$	$\hat{\alpha}$	$\hat{\beta}$	$\hat{\rho}$	$\hat{\omega}_1^*$	$\hat{\omega}_2^*$	$ln(L)^*$	$ln(L)^*$
A (IS)	0.689	0.115	5.411	0.689	0.120	0.000061	12.0	19.1	-2288.0	-2276.4
A (OOS)	0.860	0.191	2.008	0.686	0.120	0.000055	12.3	20.0	-2447.6	-2276.7
C (IS)	0.632	0.107	2.008	0.643	0.112	0.000055	17.6	34.9	-1603.7	-1596.1
C (OOS)	0.551	0.101	5.411	0.648	0.113	0.000061	17.2	33.5	-1616.1	-1596.4

Table 2: *Estimation Results–QRE vs. Rational QRE.* This table gives the estimation results from fitting Logit QRE and Probit Rational QRE to games A and C via maximum likelihood. Results are presented both in-sample ("IS") and out-of-sample ("OOS").

The first thing to note is that both QRE and Rational QRE are 1-parameter models. In the case of QRE, it is the noise parameter λ and for Rational QRE, it is the information costs parameter ρ. Since both models have the same number of degrees of freedom, their performances are directly comparable using their associated maximized likelihoods. Results are presented both *in-sample* and *out-of-sample*. In-sample implies that the parameter is estimated for game $i \in \{A, C\}$ using the data of i. Out-of-sample implies the parameter estimated for game i is used to make predictions for the play of game $j \neq i$.

The results suggest that Rational QRE outperforms QRE in-sample for both games by a small margin. What is striking is that Rational QRE significantly outperforms QRE out-of-sample for both games. In fact, Rational QRE does better *out-of-sample* than QRE does *in-sample*. I interpret these results as strong support for the Rational QRE theory. It implies that ρ is much more stable across games ordered by scale than is λ, as can be seen by comparing the estimated in-sample $\hat{\rho}$'s and $\hat{\lambda}$'s across the two games.

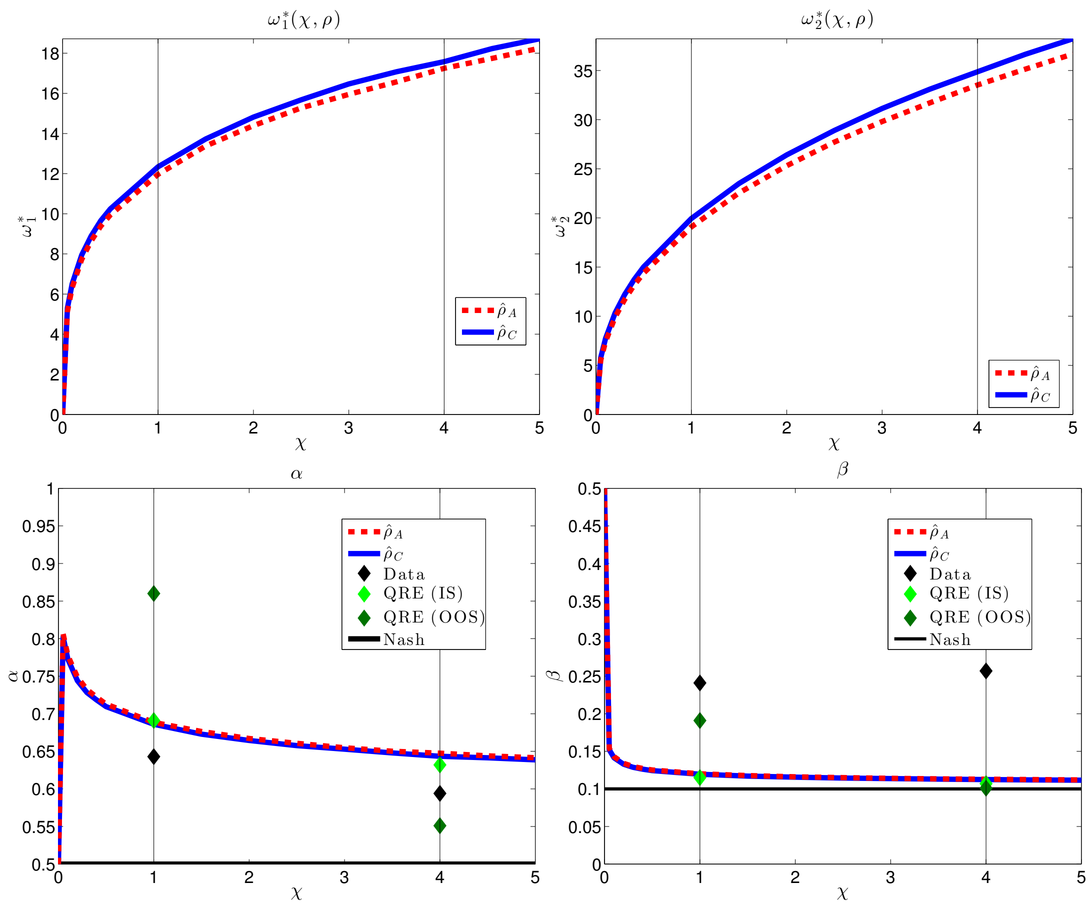

Figure 1: χ-*Correspondences for the Scale-family of* A *and* C. This figure gives predicted equilibrium precision (ω_1^*, ω_2^*) and play (α, β) as a function of scale χ for games in the scale-family of A and C. $\chi = 1$ corresponds to A and $\chi = 4$ corresponds to C. $\hat{\rho}_A$ and $\hat{\rho}_C$ are the best-fit parameters for A and C respectively.

Figure 1 depicts the Rational QRE "χ-correspondences" of games in the scale family of A, which give equilibrium objects as a function of χ. The correspondences are drawn for $\hat{\rho}_A$ and $\hat{\rho}_C$, the best-fit estimates of ρ for games A and C respectively. $\chi = 1$ and $\chi = 4$ correspond to A and C, so the couples $(\chi = 1, \hat{\rho}_A)$ and $(\chi = 4, \hat{\rho}_C)$ give in-sample estimates for A and C; and $(\chi = 1, \hat{\rho}_C)$ and $(\chi = 4, \hat{\rho}_A)$ give out-of-sample predictions for A and C. In addition, the figure is superimposed by the in-sample and out-of-sample Logit QRE predictions as well as the empirical mixing probabilities and Nash equilibrium predictions ($(\alpha, \beta) = (0.5, 0.1)$ for all χ).

From Figure 1, it is easy to see visually the strength of the out-of-sample performance: the χ-correspondences for both values of $\hat{\rho}$ nearly coincide, so either one will give similar predictions *for all* χ. Note also that for either value of $\hat{\rho}$, *attention* ω_i^* increases with χ, consistent with Conjecture 1, and play approaches Nash, consistent with Theorem 2. Of course, this would have been true for any given ρ. A better test is comparing the in-sample predictions for $\chi = 1$ and $\chi = 4$, since these are calculated independently. From this excercise, we see that subjects do in fact *seem to be* paying more attention as the scale increases, resulting in play that is more similar to Nash, and this is *not* by construction.

4 Experimental Proposal

I propose replicating the same basic experiment of McKelvey et al. [2000] in which they have subjects play game A and game C (and others) with a number of methodological differences. To summarize, their experiment is played in the lab, and each subject plays 50 rounds of one game followed by 50 rounds of another. Each session consists of 2 games, so each subject plays 100 rounds in total. Each subject is assigned to be a row player or a column player, and plays as that type for the duration of the session against randomly assigned opponents of the other type. The data used in Section 3 is the pooled data for games A and C across all sessions–even those in which only one of the two games appear.

My primary concern with the procedure is that a sizeable fraction (40%) of the data for games A and C are from subjects who play 50 rounds of A followed by 50 rounds of C or vice versa. I expect that most subjects in those sessions realized the games were ordered by scale and thus were very likely to have relied on some heuristic for adjusting their play in the second game based on the outcome of the first game (such as continuing to use the same mixed strategy). Basically, the experiment seems to have stripped subjects of the need to respond to changing stakes.

Another concern regards learning. The premise of the Rational QRE theory is that subjects make a single investment in attention per game (or round), and that investments do not carry over across games. However, since subjects play the same game many times, it is likely that attention accumulates with each round. Even if subjects are seeing the game itself as if for the first time each round, and they optimally choose attention each round, it is likely that increasing comfort with the environment makes acquiring information less costly each round. That is, under this interpretation, the information costs parameter should decrease with each round and players should act as though they are paying closer attention.

A final, minor concern is that when comparing games A and C, it is not merely the stakes that have changed, but the display observed by subjects also changes. The exchange rate is held fixed, and the payoffs are multiplied by 4 in going from A to C. If the larger numbers in the payoff matrix are more difficult for subjects to process, it may be that the assumption of a constant information costs parameter across games is violated.

Online Experiment. To circumvent these issues, I would like to run an experiment in which no individual subject plays the same game twice or any two games ordered by scale. This will prevent learning and ensure that each subject must make investments in attention freshly for each game. In addition, for games ordered by scale, I will vary the exchange rate instead of the payoff matrix. This will ensure that the complexity is held fixed across games, giving the model the best chance of being able to make strong out-of-sample predictions. The problem with this design is that it will require many more subjects for an adequate sample size. The showup fees alone would make the experiment infeasibly costly. For this reason, I propose to use an online platform such as Amazon's Mechanical Turk, where the showup fees and payments can be made much smaller, and the simplicity of the design should make it easy to implement.

5 Conclusion

QRE is a simple and elegant model that explains much of what we observe in the lab. It is flexible, powerful, and more often than not, seems to do an excellent job when compared to other models. Undoubtedly, much of the success of QRE is due to the fact that people are not the perfect agents of economic theory–like those used in the Nash equilibrium solution. It is indisputable that people make errors. This paper strikes a compromise between Nash and QRE by acknowledging that people do make errors, but that the errors across games should have certain regularities related to the strategic incentives of all players. To the extent the model is well-specified, maximum likelihood should give estimates of the information costs parameter that are both accurate and stable across games. This is what gives the model the ability to make strong out-of-sample predictions, extrapolating from one game to another. I regard this as the biggest strength of the model, and the empirical application lends support to the theory.

6 Appendix

6.1 Proof of Lemma 1

Lemma 1. *(i)* $\omega_i^{(g)*}(\beta_{-i};\chi,\rho_i)$*, as defined in (1), exists and is unique for all* (β_{-i},χ,ρ_i)*; (ii)* $\omega_i^{(g)*}(\beta_{-i};\chi,\rho_i)$ *is continuous in* (β_{-i},χ,ρ_i)*, strictly increasing in* χ*, and strictly decreasing in* ρ_i*.*

Proof. (i) $\rho_i k(\omega_i)$ is continuous in $\omega_i \in [0,\infty)$, strictly increasing, and unbounded (by assumption). $\chi\sum_{k=1}^{2} Q_{ik}^{(g)}(\omega_i,\beta_{-i})u_{ik}(\beta_{-i})$ is continuous in $\omega_i \in [0,\infty)$, strictly increasing, and bounded (obvious). This establishes existence. For uniqueness, it is sufficient to show that $\chi\sum_{k=1}^{2} Q_{ik}^{(g)}(\omega_i,\beta_{-i})u_{ik}(\beta_{-i})$ is *weakly* concave (in ω_i for given β_{-i}) so that when costs are subtracted, the objective is *strictly* concave. Under **(A1)**, $g = 1$, and we have

$$\chi \sum_{k=1}^{2} Q_{ik}^{(1)}(\omega_i, \beta_{-i}) u_{ik}(\beta_{-i}) = \chi \Phi\left(\sqrt{\frac{\omega_i}{2}} \Delta_i(\beta_{-i})\right) \Delta_i(\beta_{-i}) + const.$$

Note that if β_{-i} is such that $\Delta_i(\beta_{-i}) = 0$, this is linear and the optimum is obtained uniquely at $\omega_i^* = 0$. This is intuitive: if both actions are payoff-equivalent, then why pay for better information? If β_{-i} is such that $\Delta_i(\beta_{-i}) > 0$, then we have that $\sqrt{\frac{\omega_i}{2}}\Delta_i(\beta_{-i}) \geq 0$ for all $\omega_i \in [0, \infty)$ and thus $\chi\Phi\left(\sqrt{\frac{\omega_i}{2}}\Delta_i(\beta_{-i})\right)$ is *strictly concave* (from properties of Φ). If β_{-i} is such that $\Delta_i(\beta_{-i}) < 0$, then we have that $\sqrt{\frac{\omega_i}{2}}\Delta_i(\beta_{-i}) \leq 0$ for all $\omega_i \in [0, \infty)$ and thus $\Phi\left(\sqrt{\frac{\omega_i}{2}}\Delta_i(\beta_{-i})\right)$ is *strictly convex*, and $\chi\Phi\left(\sqrt{\frac{\omega_i}{2}}\Delta_i(\beta_{-i})\right)\Delta_i(\beta_{-i})$ is once again *strictly concave.*

Under **(A2)**, $g = 2$, and we have

$$\chi \sum_{k=1}^{2} Q_{ik}^{(2)}(\omega_i, \beta_{-i}) u_{ik}(\beta_{-i}) = \chi \frac{u_{i1}(\beta_{-i}) e^{\omega_i u_{i1}(\beta_{-i})} + u_{i2}(\beta_{-i}) e^{\omega_i u_{i2}(\beta_{-i})}}{\sum_{k=1}^{2} e^{\omega_i u_{ik}(\beta_{-i})}}.$$

This is differentiable, so is (strictly) *concave* if the second derivative is (strictly) negative, which is given by

$$-\chi \frac{(u_{i1}(\beta_{-i}) - u_{i2}(\beta_{-i}))^3 (e^{\omega_i u_{i1}(\beta_{-i})} - e^{\omega_i u_{i2}(\beta_{-i})}) e^{\omega_i (u_{i1}(\beta_{-i}) + u_{i2}(\beta_{-i}))}}{(e^{\omega_i u_{i1}(\beta_{-i})} + e^{\omega_i u_{i2}(\beta_{-i})})^3}$$

which equals 0 if $u_{i1}(\beta_{-i}) = u_{i2}(\beta_{-i})$ and strictly negative otherwise.

(ii) To show that $\omega_i^{(g)*}(\beta_{-i}; \chi, \rho_i)$ is continuous in $(\beta_{-i}, \chi, \rho_i)$, use Berge's theorem of the maximum. It is obvious that the objective is jointly continuous in $(\beta_{-i}, \chi, \rho_i)$. To invoke the theorem, all that is required in addition is a compact constraint set, but this requires some preparation as the constraint set $\omega_i \in [0, \infty)$ is not compact. Note that for any $(\hat{\beta}_{-i}, \hat{\chi}, \hat{\rho}_i)$ and any $\epsilon > 0$ there exists a $\bar{\omega}_i(\epsilon)$ such that restricting the constraint to $\omega_i \in [0, \bar{\omega}_i(\epsilon)]$ does not change the solution for all $(\beta_{-i}, \chi, \rho_i)$ within an ϵ-ball of $(\hat{\beta}_{-i}, \hat{\chi}, \hat{\rho}_i)$, and thus $\omega_i^{(g)*}(\beta_{-i}; \chi, \rho_i)$ is upper hemi-continuous for all $(\beta_{-i}, \chi, \rho_i) \in B_\epsilon(\hat{\beta}_{-i}, \hat{\chi}, \hat{\rho}_i)$ by Berge's theorem. But since $(\hat{\beta}_{-i}, \hat{\chi}, \hat{\rho}_i)$ and ϵ were chosen arbitrarily, $\omega_i^{(g)*}(\beta_{-i}; \chi, \rho_i)$ is upper hemi-continuous for all $(\beta_{-i}, \chi, \rho_i)$. And since $\omega_i^{(g)*}(\beta_{-i}; \chi, \rho_i)$ is single-valued (part *(i)*), it is also continuous.

Furthermore, that $\omega_i^{(g)*}(\beta_{-i}; \chi, \rho_i)$ is increasing in χ and decreasing in ρ_i is obvious. □

References

Andrew Caplin and Mark Dean. Revealed preference, rational inattention and costly information acquisition. *American Economic Review*, 2015. 2.5

Daniel McFadden. Quantal choice analysis: a survey. *Annals of Economic and Social Measurement*, 1976. 2.6

Richard McKelvey and Thomas Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 1995. 1

Richard McKelvey, Thomas Palfrey, and Roberto Weber. Endogenous rationality equilibrium. *Working Paper*, 1997. 1

Richard McKelvey, Thomas Palfrey, and Roberto Weber. The effects of payoff magnitude and heterogeneity on behavior in 2x2 games with unique mixed strategy equilibria. *Journal of Economic Behavior and Organization*, 2000. 1, 3, 1, 4