Implications of Quantal Response Statistical Equilibrium

Ellis Scharfenaker

Implications of Quantal Response Statistical Equilibrium

Ellis Scharfenaker
Department of Economics, University of Utah
ellis.scharfenaker@economics.utah.edu

Abstract

This paper explores the foundations and properties of the quantal response statistical equilibrium (QRSE) model developed by Scharfenaker and Foley (2017). The QRSE model provides a behavioral foundation for the formation of aggregate economic outcomes in social systems characterized by negative feedbacks. It can approximate a wide range of commonly encountered theoretical distributions that have been identified as economic statistical equilibrium and displays qualitatively similar behavior to the Subbotin and Asymmetric Subbotin distributions that range from the Laplace to the Normal distribution in the limit. Asymmetry in the frequency distributions of economic outcomes arises from the unfulfilled expectations of entropy- constrained decision makers. This paper demonstrates the logic of the QRSE model in an application to US stock market data dating back to 1926. The model provides a parsimonious explanation for the distribution of rates of return on private equities as well as clear behavioral foundation for asset price fluctuations.

1 Introduction

Many economic and social outcomes are the result of the actions of individuals interacting in an institutional structure. A significant share of economic theory tends to abstract away from the complex ways individual actions shape social outcomes by assuming that the individual actions have a zero impact on the outcome. For example, a theoretical cornerstone of Walrasian general equilibrium theory is that no individual firm or consumer can impact the equilibrium price system through their own buying and selling decisions. There are at least two obvious reasons for abstracting away from the ways individual actions shape social outcomes. First, is that the level of detail required to model and predict such a system becomes formidable when the number of individual participants becomes large. Second, information or data about individual behavior and actions tends to be either of a limited nature or is in principle unobservable.

Situations in which the actions of individual participants have a non-zero impact on the social outcome tend to produce important qualitatively different equilibrium properties of the system under analysis. In particular, such systems tend to produce central tendencies in the distribution of the outcome as well as endogenous fluctuations around these tendencies that reveal information about the underlying forces that shape the system frequency distributions. This paper examines economic systems characterized by social interactions that are conditionally dependent on perceived social outcomes but that also shape those social outcomes through stabilizing and equilibrating effects.

The quantal response statistical equilibrium (QRSE) model developed by Scharfenaker and Foley (2017) applies maximum entropy reasoning for predicting the behavior of such systems when data provides only limited information about the system.

When observations or statistical properties of social outcomes are available, but there is no direct information about the actions and behavior of individual participants that determine that social outcome, the problem of inferring the joint state space over the actions and outcome and thus predicting the configuration of the system is an ill-posed underdetermined problem.

While the unobservable component in such underdetermined models may in principle be observable, data rarely provides such information beyond the individual case-study level. Situations in which the macroeconomic outcomes are observable, but the details about individual actions that determine these outcomes are unobservable arise frequently in economics. For example, the buying/selling decisions of individual investors are determined by and determine observable equity prices, but are themselves unobservable at the system level. Firm decisions to allocate capital in particular submarkets depend on the expected profit rate and determine the observable profit rate, but only aggregate firm balance sheet data is observable. Individual labor-market decisions depend on expected remuneration and shape the income distribution, but only income frequencies are observable.

In order to make inferences about the mutual dependence of individual actions and the resulting social outcomes requires we know the structure of a joint frequency distribution over the actions and outcomes. In most economic applications both dimensions of the problem are assumed to be known (e.g. Manski and McFadden (1981)) and the problem lies only in estimating the probability of individual actions through a parametric statistical model. If either dimension of the problem is unknown, however, we face the inverse problem of having to reconstruct the joint distribution based on whatever information we have about the system. Operationally, this requires that we impose constraints on the system in the form of a prior hypothesis. One of the most successful approaches to solving such underdetermined inverse problems is the

3

Principle of Maximum Entropy (PME) inference championed by E.T. (Jaynes, 2003)[1]. The PME provides a straightforward and logical way to incorporate prior hypothesis in a model that produces the least biased estimate of such a probability distribution. A maximum entropy probability distribution is also a predictive statistical model that indicates whether or not the theoretical priors put into the calculation are relevant in the sense of providing a good fit to the data.

This paper details the logic of maximum entropy inference for underdetermined social interactions problems and examines some of the implications for the resulting statistical model. It turns out that many of the widely used statistical models used for modeling the frequency distributions of economic and social outcomes, including the Subbotin and Asymmetric Subbotin Distribution, can be usefully understood as special cases of the more general behaviorally founded QRSE model. Thus, many social phenomena that give rise to macroscopic statistical regularities turn out to have a clear interpretation from the QRSE framework.

This paper also illustrates the logic of QRSE by extending the model to monthly firm-level equity price data dating back to 1926. The results provide a remarkable fit to the data as well as a clear behavioral interpretation of asset price fluctuations that may serve as an indicator of speculative activity.

2 Statistical Equilibrium

Modeling and predicting the properties and behavior of a social system comprised of many individuals is a forbidding task. Physicists faced similar problems in the 19th century while trying to derive the phenomenal laws of gas behavior from the hypothesis that gases were constituted from particles such as molecules and atoms.

---

[1]The methods of maximum entropy inference are extensively detailed in Golan (2018).

In both cases the number of degrees of freedom in the system poses a serious problem when viewed from a deterministic predictive theory. Enumerating the details of the system is combinatorially impractical. The statistical equilibrium methods developed by Boltzmann (1871), Maxwell (1860) and Gibbs (1902) overcame this difficulty ingeniously by borrowing statistical methods from social sciences and by substituting a statistical model of the system for a detailed dynamic description of all degrees of freedom in the system.[2] The physicist E. T. Jaynes (1957) showed that the reasoning supporting statistical equilibrium methods in physics was in principle applicable to any complex system with many degrees of freedom, including biological and social systems. Jaynes' insight was that the most likely state of a system is maximally random given whatever constraints the properties of the system and its constituents put on its states.

The statistical equilibrium of a system is the probability distribution over the state space that maximizes entropy subject to whatever constraints describe the behavior of the system. The nontrivial part of modeling with statistical equilibrium methods is typically to find the constraints that parsimoniously express the theory relevant to the problem. Constraints are anything that modify the predictive statistical model. In the context of social sciences, constraints usually arise as physical constrains such as limited resources, monetary constraints such as a budget constraint, non-negativity of prices, or accounting identities, purposive human behavior such payoff maximization or "satisficing" behavior, and institutional structures such as social norms and laws. Constraints reflect the underlying theory and like any acceptable method of analysis, the predictive relevance of maximum entropy inference is conditional on the ability of the statistical model to produce observable regularities in the system under analysis. When the resulting model fails to produce empirical regularities this is an indication

---

[2]See Porter (1986) for a thorough history of statistical reasoning in the physical sciences.

that the theory and constraints are incomplete or incorrect. Many of the standard econometric principles of inference, such as likelihood analysis, hypothesis testing, and model selection follow directly from the principle of maximum entropy (Golan, 2018).

3   Quantal Responses in Social Interactions

In the simplest scenario there are just two components that determine the constraints on the system that define a statistical equilibrium in social interactions. First, individual agents respond to some economic variable, such as individual buy/sell decisions in response to price spreads in equity markets. Second, the economic variable is conditionally dependent on the actions taken, e.g. selling (buying) decision slow down (speed up) the rate of return on equities. Such statistical equilibrium models require that we think in terms of a joint frequency distribution over the actions of individuals, and the distribution of the economic variable. Letting the economic variable be $x \in \mathbb{R}$ and the action variable be $a \in \mathcal{A}$, where $\mathcal{A}$ is the set of possible actions, we are interested in the equilibrium joint distribution $f[a, x] = f[a|x]f[x] = f[x|a]f[a]$.

The first component of the model represents a behavioral theory of the typical agent in terms of the probability that she will choose a particular action a conditional on the variable x and is expressed in the conditional distribution $f[a|x]$. This conditional distribution expresses the response probabilities over the action set given x and quantifies the impact of x on the individual action frequencies. The second component of the model reflects the impact of the action on the outcome variable which is expressed in the conditional distribution $f[x|a]$. This conditional distribution expresses a theory of the formation of social outcomes. If the actions of individual participants had no impact on the outcome then $f[x|a] = f[x]$ as is the standard general equi-

librium framework. A statistical equilibrium model of the joint distribution $f[a, x]$ requires theoretical closures on the microeconomic behavioral and macroeconomic outcome dimensions of the problem.

4   Feedback Constraints

If there is a non-zero impact of individual actions on the social outcome then $f[x|a] \neq f[x]$. How individual actions shape social outcomes will depend on whether or not a particular action positively or negatively impacts x. Economic variables that exhibit statistical regularities such as a central moment and stable endogenous fluctuations can be considered in equilibrium precisely because there are negative forces that stabilize the distribution. Capitalist firms' profit rates, for example, display remarkably stable statistical properties such as a sharp peak and fat tails because the competitive behavior of firms imply that firms will tend to enter any particular market with above average profitability and exit markets with below average profitability. The act of entering, however, increases the supply in that market which drives down prices and profit rates. Exiting has the converse effect of lowering supply and rising prices and profit rates for firms that remain (Smith, 1776 [1982]). Implicit in this reasoning is the corse-graining of markets into various sub-markets in which firms compete. These might be specific goods such as watches or computers, or sub-industries such as hospitality and fast food. The negative feedback of a on x is because a affects a specific sub-market due to crowding effects and other competitive interactions.

In the simplest case the action set is the binary set $a = \{a_0, a_1\}$, for example, investors' decision to buy/sell assets in a particular sub-market, or firms' decisions to enter/exit a particular sub-market. In general, a_0 represents the "positive" action that raises x when it is below some value ξ and a_1 represents the "negative" action

that reduces x when it is above ξ. Statistical equilibrium in the joint distribution $f[a, x]$ implies that x is statistically regulated by a through negative feedbacks. The variable ξ represents the central tendency around which x statistically fluctuates. This information can be specified as a moment inequality:

$$\int_x f[a_0, x] x \, dx \leq \xi \leq \int_x f[a_1, x] x \, dx \tag{1}$$

This inequality says that the expectation of the outcome conditional on the "negative" action should be above the expectation of the outcome conditional on the "positive" action and implies the following:

$$\int_x f[x, a_1](x - \xi) \, dx - \int_x f[x, a_0](x - \xi) \, dx$$
$$= f[a_1]\mathbb{E}[(x - \xi)|a_1] - f[a_0]\mathbb{E}[(x - \xi)|a_0] \, dx \tag{2}$$
$$= \int_x (f[a_1|x] - f[a_0|x]) f[x](x - \xi) \, dx \geq 0$$

If there was no impact of the action on the outcome then maximizing the entropy of $f[a, x]$ would make this difference as large as possible. If there is a negative feedback of the action on the outcome then this difference will be constrained to some positive but finite value ϵ:

$$\int_x (f[a_1|x] - f[a_0]) f[x](x - \xi) \, dx \leq \epsilon \tag{3}$$

When ϵ is constrained to be finite agents' actions will stabilize the marginal outcome frequency, by tending to lower the outcomes for the agents who take the action a_1 and raise the outcomes for the agents who take the action a_0. Such negative

feedbacks are characteristic of a wide range of social interactions.

In Eq. 3 ϵ is an indirect measure of the dependence of a on x. The smaller ϵ is, the more effective are the actions in changing the outcome x. When $\epsilon = 0$ $\int_x f[a_0, x]x\, dx = \int_x f[a_1, x]x\, dx = \xi$ which implies the actions have an infinite effect on stabilizing the outcome at the central tendency. A very large ϵ represents a very weak impact of the action on the outcome. This interaction between the actions and outcomes is presented graphically in Figure 4 below.

Constraining the actions to impact the outcome variable in this way in addition to the normalization constraint define the maximum entropy program:

$$\underset{\{f[a,x]\geq 0}{\text{Max}} \mathcal{H}[a, x] = -\sum_a \int_x f[a, x]\log[f[a, x]]\, dx \tag{4a}$$

$$\text{subject to } \sum_a \int_x f[a, x]\, dx = 1 \tag{4b}$$

$$\text{and } \int_x (f[a_1|x] - f[a_0|x])\, f[x](x - \xi)\, dx \leq \epsilon \tag{4c}$$

To solve this maximum entropy problem, it is convenient write the joint entropy as the entropy of the marginal distribution plus the average entropy of the conditional distribution and solve for $f[x]$:

$$\mathcal{H}[a, x] = \mathcal{H}[x] + \int_x f[x]\mathcal{H}[f[a|x]]dx \tag{5}$$

$$= -\int_x f[x]\log[f[x]]dx - \int_x f[x]\sum_a f[a|x]\log[f[a|x]]dx \tag{6}$$

Thus, 4 can be equivalently written as:

$$\underset{\{f[x]\geq 0\}}{\text{Max}} \; \mathcal{H}[x] + \int_x f[x]\mathcal{H}[f[a|x]]dx \tag{7a}$$

$$\text{subject to} \; \int_x f[x]\,dx = 1 \tag{7b}$$

$$\text{and} \; \int_x \left(f[a_1|x] - f[a_0|x]\right) f[x](x-\xi)\,dx \leq \epsilon \tag{7c}$$

This programming problem has an associated Lagrangian:

$$\mathbb{L}[f[x], \lambda, \gamma, \xi] = \mathcal{H}[x] + \int_x f[x]\mathcal{H}[f[a|x]]dx - \lambda\left(\int_x f[x]\,dx - 1\right) \tag{8}$$

$$-\gamma\left(\int_x \left(f[a_1|x] - f[a_0|x]\right) f[x](x-\xi)\,dx - \epsilon\right) \tag{9}$$

Because the entropy is a strictly concave function of the probabilities, and the constraints are linear or concave functions of the probabilities, the first-order conditions are sufficient to characterize the solution:

$$\hat{f}[x] = e^{-(1+\lambda)}e^{\mathcal{H}[f[a|x]]}e^{-\gamma(f[a_1|x]-f[a_0|x])(x-\xi)} \tag{10}$$

Plugging 10 into the normalization constraint 7b and solving for $\hat{f}[x]$ gives:

$$\hat{f}[x] = \frac{1}{\mathbb{Z}[x,\gamma]}e^{\mathcal{H}[f[a|x]]}e^{-\gamma(f[a_1|x]-f[a_0|x])(x-\xi)} \tag{11}$$

where $\mathbb{Z}[x,\gamma] = \int_x e^{\mathcal{H}[f[a|x]]}e^{-\gamma(f[a_1|x]-f[a_0|x])(x-\xi)}\,dx$ is the normalizing factor or "partition function" and γ is the Lagrange multiplier associated with the constraint on the

conditional dependence of a and x.

Scharfenaker and Foley (2017) refer to this solution as the quantal response statistical equilibrium (QRSE) model. Any social system in which the quantal actions of individuals shape the social outcome through a negative feedback will tend to equilibrate in the statistical sense towards 11. The QRSE model is completely characterized by the the conditional distribution $f[a|x]$ and the Lagrange multiplier γ. The distribution $f[a|x]$ represents the action of agents conditional on the aggregate variable x and is thus a representation of the response probabilities in a decision environment. Different behavioral assumptions will lead to different closures of the QRSE model.

5   Behavioral Constraints and Expectations

In many economic settings perceived deviations from the average realization of a variable are what incentivize the actions of individual agents (Dos Santos, 2017). For example, in Scharfenaker and Foley (2017) the deviation of profit rates from the average rate of profit motivates firms to relocate capital, away from below average profit sectors and into above average profit sectors. In Yang (2018) decisions of firms to adopt a new technology are conditional on the expected average cost reduction above a "premium" required for the firm to adopt the technology. In Ömer (2018) buy/sell decisions in the housing market are conditional on the changes in the deviation of actual housing prices from the expected rate of a fundamental price increase of a house. In Blackwell (2018) buy/sell decisions are conditional on bond market spreads, and in Wiener (2019) workers' labor market participation is based on differences between a segmental wage and a reference wage.

In general, these circumstances imply the typical agents' response probability $f[a|x]$ will depend on the payoff v for choosing an action, which is the difference

11

between the expected outcome variable x and the agent's expected average payoff, or fundamental valuation of x, which we call μ such that $v[a,x] = x - \mu$. In this case, the conditional distribution will be a functional $\Omega_{f[a|x-\mu]} : v \to [0,1]$ that maps the payoff function into a response probability. Because the QRSE model captures both the micro behavioral dimension as well as the macro social outcome dimension, this specification of the payoff function v requires we differentiate between agents' expectations of the fundamental μ and the actual central tendency of the outcome, ξ.

Agents' actions are determined by the perception of the social outcome through their estimated "fundamental" value μ. Because their actions determine the social outcome x, if expectations are self fulfilling then the 3 implies $\xi = \mu$. Agents' expectations of the fundamental μ, however, need not be correct. If $\xi \neq \mu$ expectations may be unfulfilled incentivizing agents to revalue their estimate of the fundamental through some type of market-based punishment for acting incongruous with the market. If $\xi = \mu$ then expectations are fulfilled and there are no market-based incentives for the typical agent to change her estimate of μ (Foley, 2017). The difference $\zeta = \mu - \xi$ measures of the spread of the typical agent's expectations about the reference rate from the actual location of market and thus serves as a measure of unfulfilled expectations. In some contexts this divergence of expectations from the actual market can be usefully interpreted as a measure of speculative activity. Accordingly, we can write the solution to 11 as:

$$\hat{f}[x] \propto e^{\mathcal{H}[f[a|x]]} e^{-\gamma\left(\Omega_{f[a_1|x-\mu]} - \Omega_{f[a_0|x-\mu]}\right)(x-\xi)} \tag{12}$$

Individual expectations shape the outcome only to the extent that $\Omega_{f[a|x-\mu]}$ influ-

ences $\hat{f}[x]$. Only in the special case when ξ is constrained to equal μ will individual agents be able to fully determine the reference locations of the outcome through expectations alone, i.e. expectations are rational and self fulfilling in line with Friedman (1957) and Muth (1961).

5.1 Logit Quantal Responses

Research dating back to the work of Gustav Fechner and Ernst Weber has long recognized the limited cognitive capacity of individual decision makers leading to the observation that individual choice behavior is not always optimal. In economics, Herbert Simon (1955) introduced an analog to non-optimizing choice behavior in what he called "bounded rationality." Interpreted from the perspective of the theory of individual choice behavior, bounded rationality implies that choice decisions are best described as a probabilistic phenomenon as opposed to the deterministic rational theory of choice which assumes choices are always associated with probabilities equal to unity.

The probabilistic approach to individual choice was most famously introduced by the psychologist Duncan Luce (Luce & Raiffa, 1957; Luce, 1959; Suppes, Krantz, Luce, & Tversky, 1989) and was developed along game theoretic and econometric lines by McFadden (1976, 2001), McKelvey and Palfrey (1995, 1998), and Manski and McFadden (1981). Both Luce's choice axioms and the random utility models developed by McFadden lead to multinomial logit functions for discrete choice probabilities. More recently Matějka and McKay (2015), Scharfenaker and Foley (2017), Sims (2003), Wolpert (2006) have derived the same choice probabilities by identifying bounded rationality with informational processing constraints, or "entropy-constrained behavior."

If the typical agent chooses a mixed strategy, $f[a|x] : \mathcal{A} \times X \to (0,1)$ over actions $a \in \mathcal{A}$ to maximize the expected payoff, $\sum_a f[a|x]v[a,x]$ subject to a minimum informational entropy then the agent's expected payoff maximization program and associated Lagrangian are:

$$\underset{\{f[a|x]\geq 0\}}{\text{Max}} \sum_a f[a|x]v[a,x]$$

$$\text{subject to} \sum_a f[a|x] = 1$$

$$-\sum_a f[a|x]\log[f[a|x]] \geq H_{min} \tag{13}$$

$$\mathcal{L} = -\sum_a f[a|x]v[a,x] - \lambda\left(\sum_a f[a|x] - 1\right)$$

$$+ T\left(\sum_a f[a|x]\log[f[a|x]] - H_{min}\right)$$

The solution to this programming problem is the logit quantal response (LQR) distribution:

$$f[a|x] = \frac{e^{\frac{v[a,x]}{T}}}{\sum_{a'\in\mathcal{A}} e^{\frac{v[a,x]}{T}}} \tag{14}$$

The Lagrange multiplier T represents the attentiveness of the typical agent. The lower is T the more alert the actor is to differences in payoff, and the more closely the action approximates the unconstrained payoff-maximizing outcome. While entropy-constrained behavior is a generalization of rational choice theory there are several important differences with far reaching implications. When $T > 0$ individual agents preferences no longer satisfy the assumptions of consistency and completeness. Presented with the same set of options entropy-constrained individuals with identical preferences will sometimes choose one option and sometimes the other. More im-

portantly, many of equilibrium conditions and welfare implications based on uncon-
strained rational choice no longer hold when $T > 0$.

Entropy-constrained behavior also leads to the logit quantal response distribution
without having to imposing awkward distributional assumptions on the errors that
affect decision making. Rather than treating the distribution over the individual's
mixed strategy as arising from random error of noisy players or oscillating utility
as is the convention in random utility theory, the informational entropy constrained
model gives meaning to the observed dispersion of behavior as the the relative payoff
of different actions.

Alternatively, we can view the problem from the dual perspective of the observer
of the system in order to find the most probable distribution of actions. Maximizing
the entropy of the mixed strategy subject to a constraint on the expected payoff being
a minimum $\bar{v}$:

$$
\begin{aligned}
& \underset{\{f[a|x] \geq 0\}}{\text{Max}} -\sum_a f[a|x] \log[f[a|x]] \\
& \text{subject to } \sum_a f[a|x] = 1 \\
& \qquad\qquad \sum_a f[a|x] v[a, x] \geq \bar{v}
\end{aligned}
\tag{15}
$$

The associated Lagrangian to this dual problem is:

$$
\begin{aligned}
\mathcal{L} = & -\sum_a f[a|x] \log[f[a|x]] - \lambda \left(\sum_a f[a|x] - 1 \right) \\
& + \beta \left(\sum_a f[a|x] v[a, x] - \bar{v} \right)
\end{aligned}
\tag{16}
$$

In this case we are imposing a constraint on our uncertainty of the distribution of
actions conditional on the outcome for a set of agents subject to the constraint that

individuals have a minimum expected payoff for acting, which might be viewed as a "participation constraint" or reflecting what Herbert Simon referred to "satisficing" behavior, a form of bounded rationality in which agents do not maximizing their payoff *per se*, but are content with doing "well-enough" by some criterion. The most probable distribution of actions (the one with the maximum informational entropy) is again:

$$f[a|x] = \frac{e^{\beta v[a,x]}}{\sum_{a' \in \mathcal{A}} e^{\beta v[a,x]}} \tag{17}$$

For the binary action case with the linear payoff function $v[a,x] = x - \mu$ the solution becomes:

$$f[a_0|x] = \frac{1}{1 + e^{\frac{x-\mu}{T}}} \tag{18}$$

$$f[a_1|x] = \frac{1}{1 + e^{-\frac{x-\mu}{T}}} \tag{19}$$

where $T = 1/\beta$. The LQR conditional probabilities capture a wide rage of probabilistic choice behavior as is shown in Figure 1.

As long as $T > 0$ there is a positive probability for each action an individual will make. As $T \to 0$, the LQR function converges to the Heaviside step function and as $T \to \infty$ the actions of individual agents are random and conditionally independent of the outcome x and $f[a|x] \to f[a]$. These two limiting cases have important implications for the QRSE distribution.
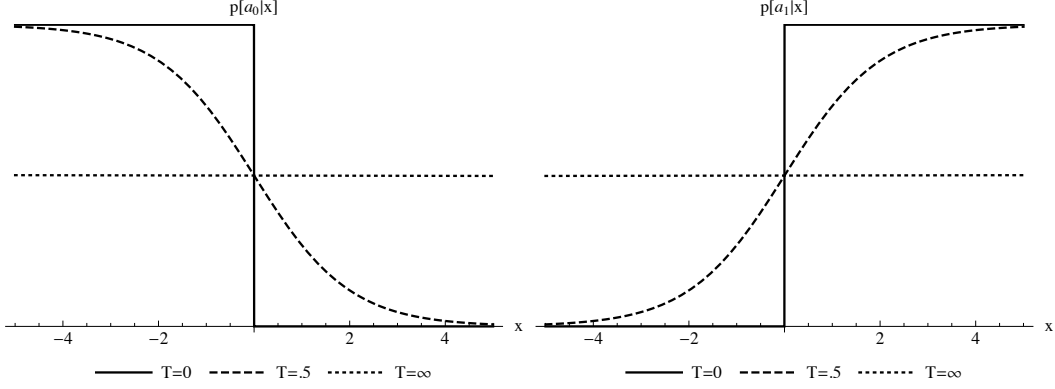
Figure 1: Logit quantal response conditional probabilities.

5.2 Fulfilled Expectations

Plugging the LQR conditional probabilities into 11 and assuming expectations are fulfilled such that $\xi = \mu$ we get the following expression:

$$\hat{f}[x] \propto \left(\frac{1}{1 + e^{-\frac{x-\mu}{T}}}^{-\frac{1}{1 + e^{-\frac{x-\mu}{T}}}} \frac{1}{1 + e^{\frac{x-\mu}{T}}}^{-\frac{1}{1 + e^{\frac{x-\mu}{T}}}} \right) e^{-\gamma \tanh\left[\frac{x-\mu}{2T}\right](x-\mu)} \tag{20}$$

$$= e^{\mathcal{H}_{\mu,T}[x]} e^{-\gamma \tanh\left[\frac{x-\mu}{2T}\right](x-\mu)} \tag{21}$$

Where $T > 0$, $\gamma > 0$, and $\mu \in \mathbb{R}$. Eq. 21 turns out to have some important properties and wide ranging implications as pointed out in Blackwell (2018).[3]

In the limiting case where $T \to 0$ agents will always chose the action a_1 whenever $x > \mu$ and a_0 whenever $x < \mu$. In this case the conditional frequencies are:

$$f[a_1|x] = \theta[x - \mu] \tag{22}$$

$$f[a_0|x] = 1 - \theta[x - \mu] \tag{23}$$

---

[3]In the spirit of the maximum entropy formalism it is worth noting that the QRSE model can be derived in a single optimization problem that includes the constraints on $f[a|x]$.

where θ is the Heaviside step function. Plugging these response probabilities into 11 we find:

$$\hat{f}[x] \propto e^{-\gamma|x-\mu|} \tag{24}$$

Solving for the constraints 7b and 7c produces $\gamma = \epsilon^{-1}$ and $\mathbb{Z}[x, \gamma] = \frac{2}{\gamma}$. Thus $\lim_{T \to 0} \hat{f}[x] = \text{Laplace}[\mu, \gamma^{-1}]$.
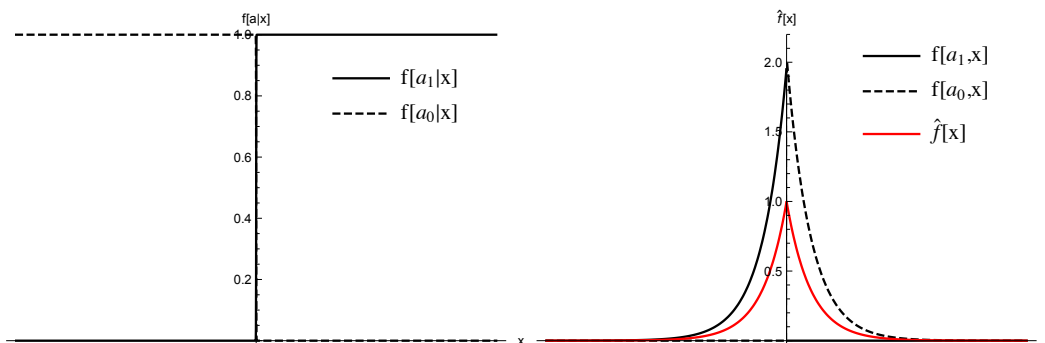

Figure 2: Marginal, conditional, and joint frequency distributions for Heaviside response probabilities for $\gamma = 2$.

Interestingly, the Heaviside response probabilities predict Laplace distributed fluctuations of the outcome variable and not a degenerate Dirac delta function $\hat{f}[x] = \delta[x - \mu]$. This result implies that the fluctuations of x are determined by the feedback constraint and not just the response probabilities. Only as the constraint becomes increasingly binding and $\gamma \to \infty$, does the the entropy $\mathcal{H}[x] = \log[\frac{2e}{\gamma}]$ decreases to its minimum and the endogenous fluctuations of x reduce to zero. Systems of social interactions producing negative feedbacks that are attended to with infinite precision thus stabilize the social outcome into a Laplace distribution. In this extreme case the partition function $\mathbb{Z}[x, \gamma]$ has a simple analytical solution. This result can act as a benchmark for social interaction models with self-fulfilling expectations and infinitely rational agents. While the Laplace distribution implies a zero-entropy constraint, it

may be difficult in many situations to distinguish between $T = 0$ and $T \approx 0$. When T is very low $f[x]$ may appear to be Laplace distributed, but as mentioned above, the same conclusions cannot be drawn from the knife-edge case when $T = 0$.

On the other hand, when $T \to \infty$ actions become conditionally independent of the outcome and the response probabilities become random. When this happens the marginal outcome frequencies $f[x]$, ceterus peribus, converge to a uniform distribution. We need to consider, however, that there is an implicit interaction between T and γ. This multivariate limit can be easily calculated by transforming the parameters into polar coordinates. Letting $\gamma = r \cos[\theta]$ and $T = r \sin[\theta]$ for $\theta \in [0, \pi/2]$

$$
\begin{aligned}
\lim_{T \to \infty, \gamma \to \infty} e^{-\gamma \tanh\left[\frac{x-\mu}{2T}\right](x-\mu)} &= \lim_{r \to \infty} e^{-r \cos[\theta] \tanh\left[\frac{x-\mu}{2r \sin[\theta]}\right](x-\mu)} \\
&= e^{-\frac{\cot[\theta](x-\mu)^2}{2}} \\
&= e^{-\frac{\gamma(x-\mu)^2}{2T}}
\end{aligned}
\tag{25}
$$

Because $f[a_1|x] - f[a_0|x] = 1/2 - 1/2 = 0$ in the limit, $\epsilon = 0$ as the logic of the constraint would confirm. Solving for the normalization constraint produces $\mathbb{Z}[x, \gamma] = \sqrt{2\pi \frac{T}{\gamma}}$ implying

$$
\lim_{T \to \infty, \gamma \to \infty} \hat{f}[x] = \mathrm{Normal}[\mu, \frac{T}{\gamma}]
\tag{26}
$$

In this case marginal, conditional, and joint frequencies are equal as is shown in Figure 3. The corresponding entropy is $\mathcal{H}[x] = \frac{1}{2} \log[2\pi e \bar{T}]$.

In between these the Laplace and Normal limiting cases the QRSE distribution with fulfilled expectations is capable of generating a wide range of symmetric distribution functions and approximating many well-known distributions with varying kurtosis much like the Subbotin distribution. Figure 4 plots the marginal and joint
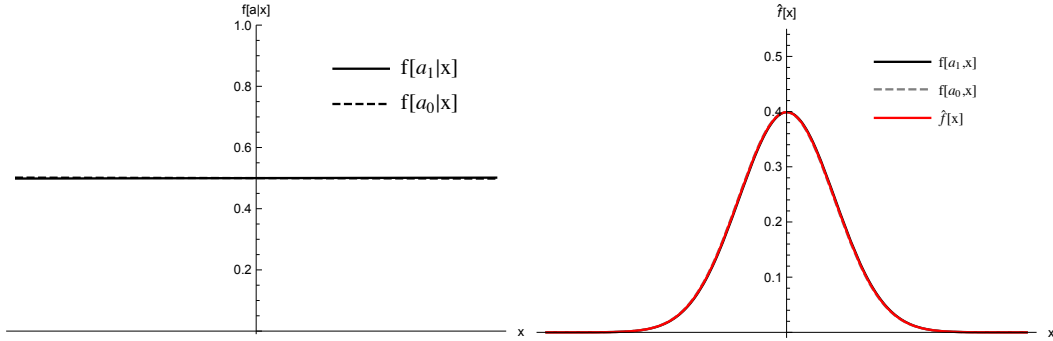
19

Figure 3: Marginal, conditional, and joint frequency distributions are all equal in the limiting case when $T \to \infty$ and $\gamma \to \infty$.

predicted frequencies from the QRSE model when agents face LQR probabilities and have fulfilled expectations for a range of parameter values.

The feedback constraint corresponding to Eq. 3 is clearly seen in middle row of Figure 4. Because γ is the Lagrange multiplier associated with the constraint it is the dual to and varies inversely with ϵ. When γ is very small the difference in the weighted expected outcomes is very large and the actions have a negligible impact on stabilizing the marginal distribution $f[x]$. When γ is very large this difference is very small and $f[x]$ becomes highly peaked. Conceptually, as $\epsilon \to \infty$, ceterus paribus, we would expect $f[x]$ to become uniform. The parameter γ, however, interacts with T in determining $f[x]$ such that when $\gamma \to 0$ the $\text{var}[x] \to \infty$ but strictly speaking $f[x]$ only becomes uniform when $T \to \infty$ as well.

The use of symmetric distributions with varying kurtosis have been identified as equilibrium distributions in a variety of economic contexts, including the distribution of firm growth rates (Bottazzi & Secchi, 2006; Stanley et al., 1996), firm profit rates (Alfarano, Milaković, Irle, & Kauschke, 2012; Scharfenaker & Semieniuk, 2016), global GDP growth rates (Williams, Baek, Li, Park, & Zhao, 2017), interest rates (Klein, 1993; Kozubowski & Podgórski, 1999), stock market returns (Madan & Seneta, 1990), option pricing (Longstaff, 1995), and currency exchange rates (Kotz,
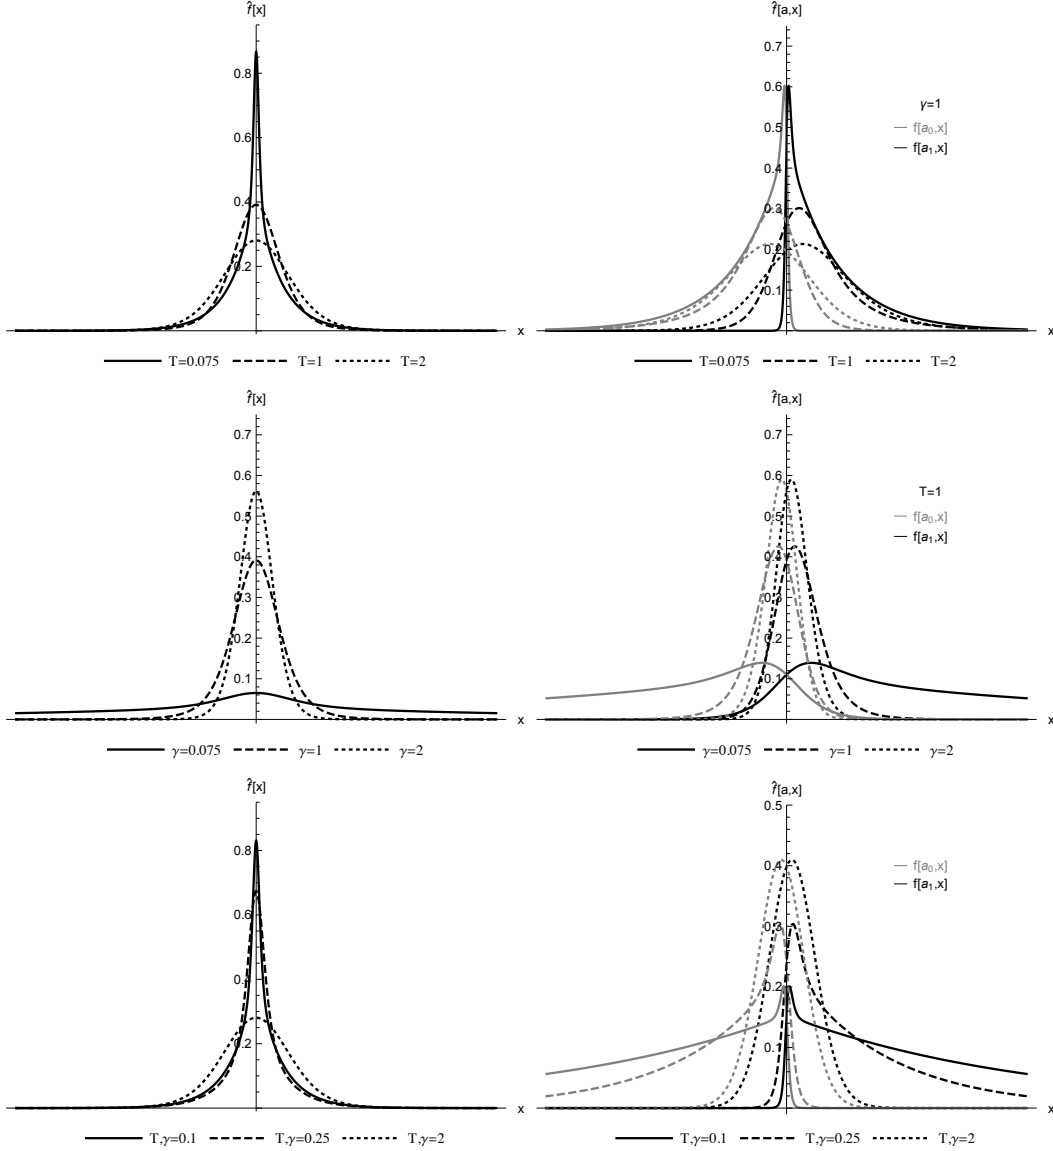
Figure 4: Marginal and joint frequency distributions for LQR probabilities. Top: $\gamma = 1$ while T varies from 0.075 to 2, Middle: $T = 1$ while γ varies from 0.075 to 2, Bottom: Varying T and γ together from 0.1 to 2.

Kozubowski, & Podgórski, 2001). In many of these models there is often an explicit if not implicit behavioral theory that determines the frequency distribution of the outcome. The QRSE model makes this dimension explicit in the model and can thus serve as an effective alternative to these equilibrium distributions. It is also a

behaviorally founded alternative to the widely used Subbotin distribution (Alfarano & Milaković, 2008; Bottazzi & Secchi, 2006; Mundt, Alfarano, & Milaković, 2015; Subbotin, 1923) for distributions with excess kurtosis between 0 and 3, for which the Subbotin produces Laplace and Normal outcomes[4].

5.3 Unfulfilled Expectations

When agents' expectations of the market are allowed to be incorrect μ is no longer constrained to be equal to ξ. If the typical agent also faces behavioral constraints in the form of the LQR conditional probabilities the QRSE model for $\mu, \xi \in \mathbb{R}$ is defined by the kernel:

$$\hat{f}[x] \propto e^{\mathcal{H}_{\mu,T}[x]} e^{-\gamma \tanh\left[\frac{x-\mu}{2T}\right](x-\xi)} \tag{27}$$

Allowing for the possibility that expectations of the average social outcome $\mu \neq \xi$ introduces a skew to the distribution. An important implication of the general QRSE model with unfulfilled expectations is that not only do individual agents behave probabilistically according to the LQR function, but agents beliefs can be mistaken leading the typical agent to, for example, buy or sell at too high or too low a rate of return relative to their expectations. When agents act on mistaken beliefs and expectations go unfulfilled the distribution of the social outcome becomes asymmetric.[5] Figure 5 plots the marginal, conditional, and joint frequency distributions for the QRSE model

---

[4]Empirically, these are the relevant limit values for applications of the Subbotin distribution. Most economic data is shown to approximate the Laplace limit. To this author's knowledge, no application of the Subbotin distribution has identified excess kurtosis greater than 3.

[5]Dos Santos and Scharfenaker (2019) come to a similar conclusion and identify unfulfilled expectation as indicative of bullish speculator behavior. Recently, Hansen, Szöke, Han, and Sargent (2019) have been developing "twisted relative entropy" probabilities to reflect such mistaken beliefs.
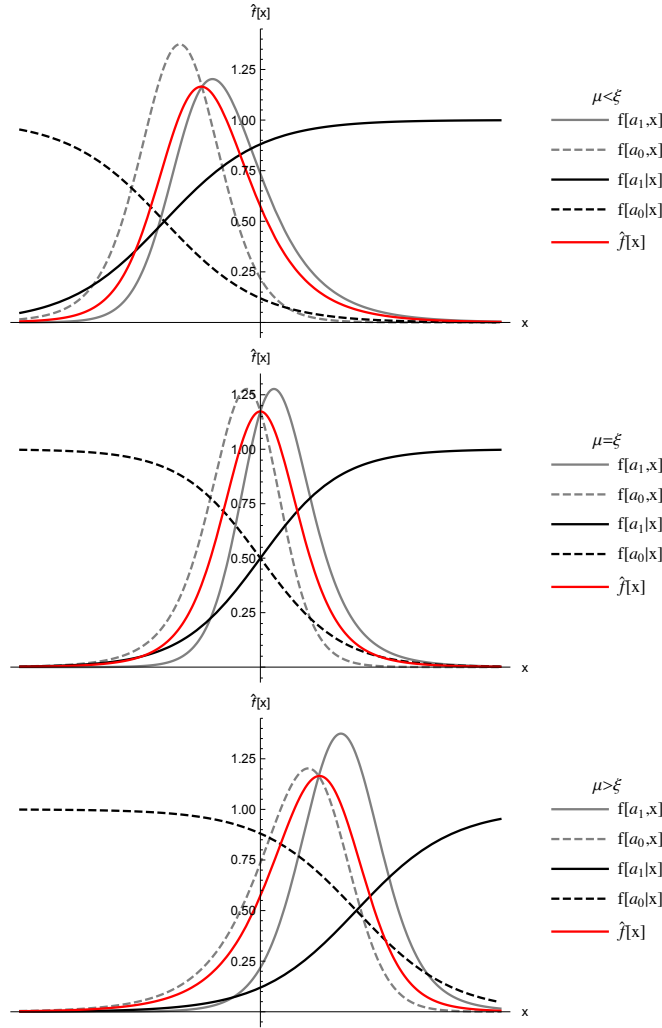
with unfulfilled expectations.

Figure 5: Marginal, conditional, and joint frequency distributions for LQR response probabilities and unfulfilled expectations. Top: $\mu = -2, \gamma = 1, T = 1, \xi = 0$, Middle: $\mu = 0, \gamma = 1, T = 1, \xi = 0$, Bottom: $\mu = 2, \gamma = 1, T = 1, \xi = 0$.

The general form of the QRSE model with mistaken beliefs also approximates many well-known closed form skew distributions and may serve as a useful alternative to the five parameter Asymmetric Power Exponential Distribution (APED) or Asymmetric Subbotin Distribution (Bottazzi & Secchi, 2011), which has been used to model the statistical equilibrium of outcomes such as the profit rate (Mundt &

Oh, 2019), consumption expenditures (Fagiolo, Alessi, Barigozzi, & Capasso, 2010), and firm growth rates (Bottazzi, Li, & Secchi, 2019). One of the major obstacles with the APED and similar distributions is justifying the modulus in the kernel of the pdf in a theoretically compelling way. As the logic of the QRSE model shows, the APED assumes that agents or firms maximize the function $v[a,x] = |\frac{x-\mu}{\sigma}|^{\alpha}$, i.e. maximize the scaled absolute value of deviations from their estimated fundamental with a constrained dispersion, which may only make economic sense in some contexts. In the profit rate example (Alfarano et al., 2012; Mundt et al., 2015; Mundt & Oh, 2019; Scharfenaker & Semieniuk, 2016) capitalist firms are not aiming at a target profit rate and disappointed equally if they exceed or fall short of it as the modulus function implies. They are simply trying to get the highest profit rate. It is the feedback of firms' entry and exit decisions that stabilizing the resulting distribution of profit rates, which is only made explicit in the QRSE model. Further, the modulus function implies the limiting Heaviside step behavior, which effectively assumes zero informational entropy on the typical agent's mixed strategy, an assumption with little empirical support.

The general QRSE model describes systems in which the economic outcome is brought into statistical equilibrium by the deliberate actions of boundedly rational agents and whose expectations may be mistaken. It can capture many of the same economic phenomena as the AEPD for $0 \leq \alpha \leq 3$, but offers the benefit of interpretability and theoretical consistency for models that implicitly or explicitly rely on behavioral foundations, but fail to incorporate this dimension as a determining factor of the economic outcome.

6 Application to US Stock Market Data

The QRSE model can be easily applied to the market for tradable securities and provide a behavioral foundation for the distribution of asset prices. Blackwell (2018), for example, develops a Keynesian model of asset pricing and applies the QRSE distribution to cross-sectional US equity returns in the Wikiprice database for daily, weekly, and monthly realized returns from 1990 to 2017. The model presented here follows the same logic, but examines a different historical dataset dating back to 1929.

In the simplest scenario, a QRSE model of asset prices implies that individuals buy and sell some asset in a large number of separate sub-markets based on the rate of increase in the asset's price. In this model individual stocks are used as proxies for the corse-grained sub-markets. The action in this model is $a = \{\text{buy, sell}\}$. The frequency of buying/selling decisions in a particular sub-market is conditional on the actual rate of return $\hat{p}$ of an asset, which is the outcome of the social interactions of all stock market participants. The expectation of the market participants μ, is the estimate of the fundamental rate of price increase of an asset. The typical agent's payoff function is the difference between the two: $v[a, \hat{p}] = \hat{p} - \mu$. Accordingly, agents' behavior is described probabilistically by the logit-quantal response function:

$$f[\text{buy}|\hat{p}] = \frac{1}{1 + e^{\frac{\hat{p}-\mu}{T}}} \tag{28}$$

$$f[\text{sell}|\hat{p}] = \frac{1}{1 + e^{-\frac{\hat{p}-\mu}{T}}} \tag{29}$$

Individual acts of buying or selling have a non-zero impact on the rate of return in the sub-market such that buying assets will raises the price faster and selling will slow it down. This impact of actions on outcomes limits the rate of price appreciation conditional on buying and selling, which constrains the equilibrium marginal

frequency distribution $f[\hat{p}]$. When the equilibrium distribution of rates of return is asymmetric, the typical agent incurs losses relative to the average market rate. These losses appear as long tails and a skew to $f[\hat{p}]$.

Using monthly CRSP price data for all exchanged securities on the US stock market from 1926-2018 Eq. 27 is fit to the average annual rate of return as measured by the log return.[6] This data comprises on average $37,990$ observations per year with a total of $3,533,145$ observations spanning over a 90 year period. Estimation is done using a Robust Adaptive Metropolis (RAM) algorithm (Vihola, 2012)[7].
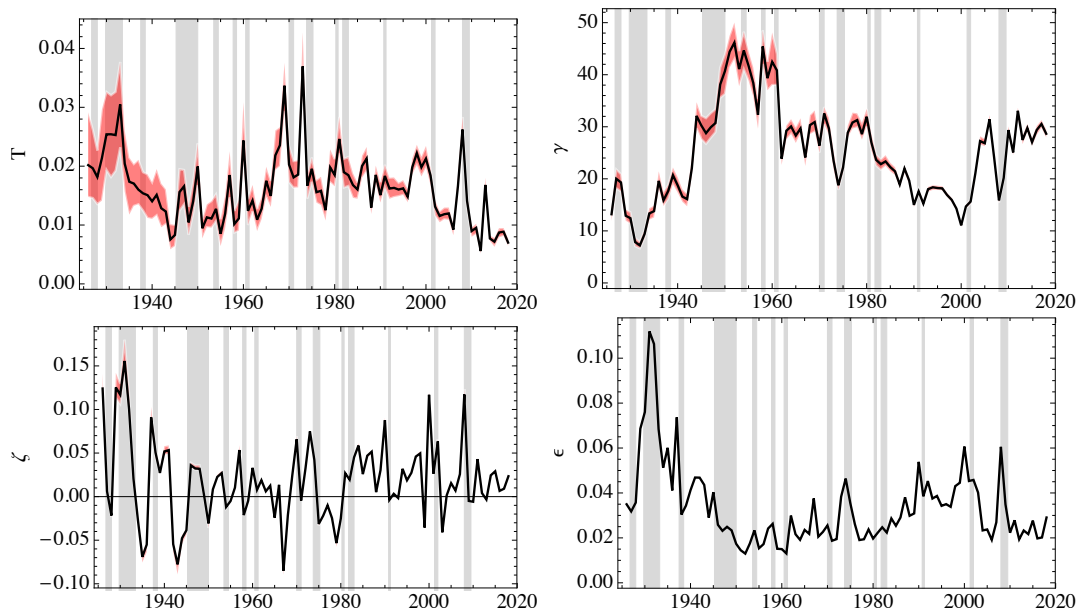

Figure 6: QRSE parameter estimates for the rate of price increases across all US security exchange markets 1929-2018 with 95% credibility intervals. NBER recession dates are plotted as gray bars. Data source CRSP.

Figure 6 plots the evolution of QRSE model parameters applied to cross sections of the rate of price increases for US securities for over 90 years. The results are telling and confirm the similar findings of Blackwell (2018). The behavioral temperature T,

---

[6]Using annual average rather than the monthly data eliminates much of the high-frequency noise and is done for visual clarity. See Appendix A: for details on data.

[7]R code is available upon request.

which represents the responsiveness of individual agents to changes in the rate of return appears in line with business cycle frequencies as well as lower frequencies. It is confined to low, but positive levels throughout the sample which implies that the typical agent's buy/sell decisions are probabilistic and not deterministic as with the Heaviside step response function. It appears that in times leading up to a recession perceptual limitations of individual investors to price differentials become increasingly binding and that after a downturn T is brought back to its trend. This result could be due to the increased amount of trading activity in booms that leads to greater inattentiveness in market participants.

The parameter γ is dual to the constraint ϵ and thus these variables move in opposition of one another. The parameter ϵ is easier to interpret, it represents the limiting difference between the feedback of buy/sell decisions on the rate of return. This spread tends to increase leading up to a downturn which implies a less binding feedback constraint and declining γ. Because the feedback constraint implies that buying decisions raise the rate of prices faster and selling decisions slow it down, the spike in ϵ prior to a downturn indicates that buying decisions in speculative times e.g. from 2001 to about 2006 tend to have a strong impact on the rate of return, but that this effect weakens leading up to a recession.

The parameter $\zeta = \mu - \xi$ measures the spread between individual's expectations about the fundamental value of rates of return and the actual average market outcome. When $\zeta > 0$ expectations of the rate of return outpace their actual value. Not surprisingly, this value tends to peak prior to a downturn. ζ thus appears a leading indicator for recessions that may be interpreted as a measure of speculative activity (Dos Santos & Scharfenaker, 2019).

For comparison, the data and the QRSE model for 2004, a year in the middle of a business cycle is compared to 2008 in Figure 7:
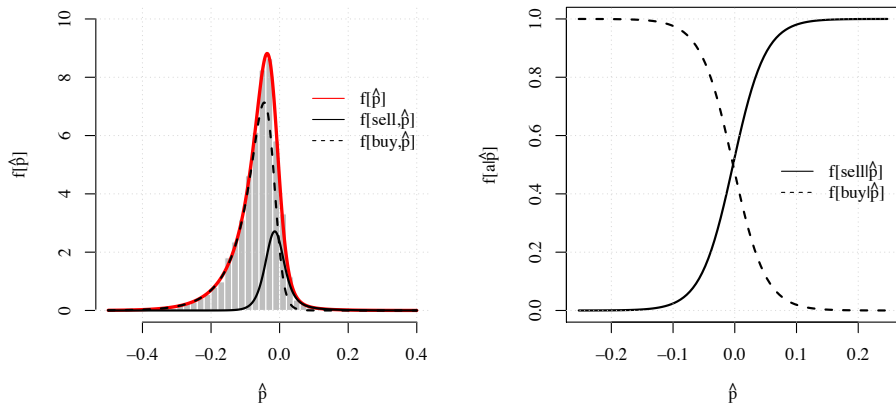
2004
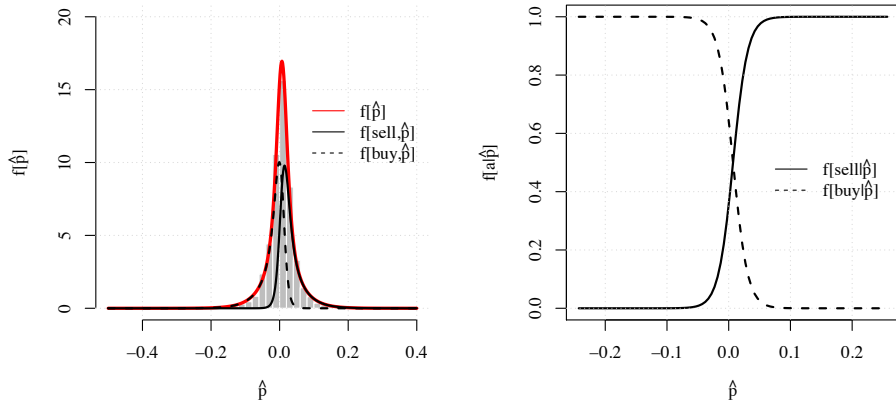
2008

Figure 7: Frequency distribution of $\hat{p}$ with QRSE model fits 2004 and 2008.

Looking at the estimated marginal frequencies $f[buy]$ and $f[sell]$ we can also see that investors' are far more likely to buy leading up to a recession. From the logic of the model this suggests that expectations of gains are due to a perception of undervaluation of asset prices. Eventually, this exuberance is brought into check by some punitive mechanism and expectations are dramatically reversed during and following a recession. The stationarity of ζ around zero is a good indication that such behavioral is cyclical and the relation to $f[a]$ suggests it is also endogenous to the system.
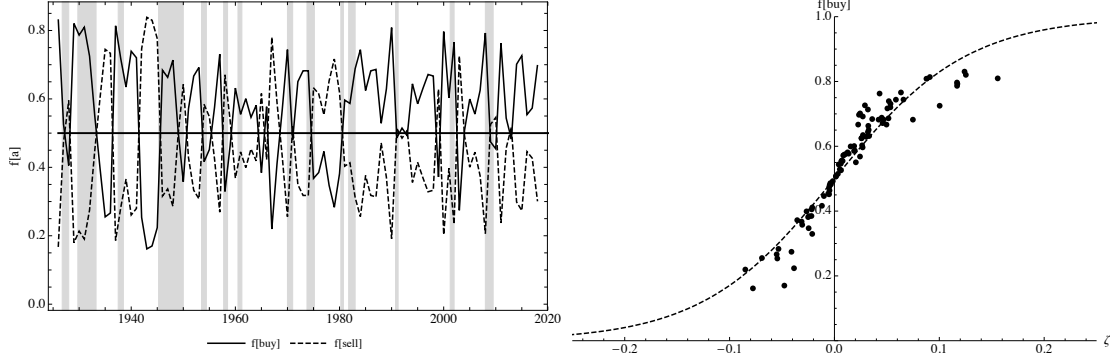
Figure 8: Left: Marginal frequencies of buying/selling 1926-2018. NBER recession dates are plotted as gray bars. Data source CRSP. Right: Relationship between buying frequencies and ζ.

7 Discussion and Conclusions

The principle of maximum entropy has been applied fruitfully to many economic and social situations. One limitation of such applications, however, is the tendency neglect of the *joint* determination of social outcomes and the actions that shape them in the formal modeling despite the fact that both dimensions of the problem are typically articulated in the theoretical exposition. The fact that closed-form marginal distributions such as the Laplace or Subbotin distributions are good model candidates for many economic variables in statistical equilibrium based on fit has the unfortunate consequence of leading to economic rationalization based on mathematical necessity of the constraints.[8] From the Principle of Maximum Entropy perspective (Jaynes, 2003) the use of closed-form distributions seems somewhat premature. While there may be good reasons to rely on such distributions for modeling statistical equilibrium in some situations, in general this limits inference to an arbitrary subset of models. The QRSE model does not have a closed-form solution, but due to the low dimensionality of the

---

[8]This reasoning is admittedly evident in Scharfenaker and Semieniuk (2016) who attempt to rationalize the Laplace and Asymmetric Laplace as the equilibrium distribution of profit rates based on entry/exit dynamics without explicitly considering the how those dynamics shape the profit rate distribution.

model, the partition function is more an inconvenience than an obstacle.

As mentioned above the difficult part of using the method of maximum entropy is to find the constraints that parsimoniously express the theory relevant to the problem. In this respect the QRSE model makes some milage over the univariate models that focus only on the marginal distribution of the social outcome by incorporating the causal mechanism as an additional dimension of the statistical problem. The QRSE model is also a very good example of the potential of maximum entropy reasoning in social sciences and may be used as a template for more complex problems that require careful reasoning in formulating the correct constraints consistent with theory.

The model developed by Scharfenaker and Foley (2017) is applicable to a wide range of social interaction problems. Any system constrained by negative feedbacks and boundedly rational individuals will tend to generate outcomes of the QRSE form. Not surprisingly, due to the limited constraints placed on the system, the QRSE model is highly flexible and capable of generating many unimodal symmetric and asymmetric distributions, including those with fat-tails. By differentiating between individual expectations and average social outcomes the QRSE model is also capable of bridging an important theoretical gap between the micro-behavioral and the macroeconomic dimensions. Because the maximum entropy framework only incorporates the information put into the calculation, each parameter in the statistical equilibrium model has a clear theoretical interpretation consistent with the constraint.

References

Alfarano, S., & Milaković, M. (2008). Does classical competition explain the statistical features of firm growth? *Economics Letters, 101*(3), 272–274.

Alfarano, S., Milaković, M., Irle, A., & Kauschke, J. (2012). A statistical equilibrium model of competitive firms. *Journal of Economic Dynamics and Control, 36*(1), 136–149.

Blackwell, K. (2018). *Entropy constrained behavior in financial markets: A quantal response statistical equilibrium approach to financial modeling.* The New School for Social Research.

Boltzmann, L. (1871). Über das wärmegleichgewicht zwischen mehratomigen gasmolekülen. *Wiener Berichte, 63,* 397–418.

Bottazzi, G., & Secchi, A. (2006). Explaining the distribution of firm growth rates. *The RAND Journal of Economics, 37*(2), 235–256.

Bottazzi, G., Li, L., & Secchi, A. (2019). Aggregate fluctuations and the distribution of firm growth rates. *Industrial and Corporate Change, 28*(3), 635–656.

Bottazzi, G., & Secchi, A. (2011). A new class of asymmetric exponential power densities with applications to economics and finance. *Industrial and Corporate Change, 20*(4), 991–1030.

Dos Santos, P. L. (2017). The principle of social scaling. *Complexity, 2017*(ID 8358909).

Dos Santos, P. L., & Scharfenaker, E. (2019). Competition, self-organization, and social scaling - accounting for the observed distributions of tobin's q. *Industrial and Corporate Change, Forthcoming.*

Fagiolo, G., Alessi, L., Barigozzi, M., & Capasso, M. (2010). On the distributional properties of household consumption expenditures: The case of italy. *Empirical Economics, 38*(3), 717–741.

Foley, D. (2017). *Information theory and behavior* (Working Paper No. 31/2017). The New School for Social Research.

Friedman, M. (1957). *A theory of the consumption function.* Princeton, NJ: Princeton University Press.

Gibbs, J. W. (1902). *Elementary principles in statistical mechanics.* New York: C. Scribner.

Golan, A. (2018). *Foundations of info-metrics: Modeling, inference and imperfect information.* New York, NY: Oxford University Press.

Hansen, L. P., Szöke, B., Han, L. S., & Sargent, T. J. (2019). *Twisted probabilities, uncertainty, and prices* (Becker Friedman Institute For Economics Working Paper No. 2019-02). University of Chicago.

Jaynes, E. T. (2003). *Probability theory: The logic of science.* Cambridge University Press.

Jaynes, E. T. (1957). Information theory and statistical mechanics. *The Physical Review, 106*(4), 620–630.

Klein, G. E. (1993). The sensitivity of cash-flow analysis to the choice of statistical model for interest rate changes. *Transactions of the Society of Actuaries, 45*, 79–186.

Kotz, S., Kozubowski, T., & Podgórski, K. (2001). *The laplace distribution and generalizations: A revisit with new applications to communications, economics, engineering, and finance.* Birkhauser Verlag GmbH.

Kozubowski, T., & Podgórski, K. (1999). A class of asymmetric distributions. *Actuarial Research Clearing House, 1*, 113–134.

Longstaff, F. A. (1995). *Stochastic volatility and option valuation: A pricing-density approach.* The Anderson Graduate School of Management, UCLA.

Luce, D. R., & Raiffa, H. (1957). *Games and decisions.* New York: Wiley.

Luce, D. R. (1959). *Individual choice behavior.* New York: Wiley.

Madan, D., & Seneta, E. (1990). The variance gamma (v.g.) model for share markets returns. *Journal of Business, 63*, 511–524.

Manski, C. F., & McFadden, D. (1981). *Structural analysis of discrete data with econometric applications*. MIT Press.

Matějka, F., & McKay, A. (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review, 105*(1), 272–298.

Maxwell, J. C. (1860). Illustrations of the dynamical theory of gases. *Philosophical Magazine, 4*(19).

McFadden, D. (1976). Quantal choice analysis: A survey. *Annals Of Eocnomic and Social Measurement, 5*(4), 363–390.

McFadden, D. (2001). Economic choices. *The American Economic Review, 91*(3), 351–378.

McKelvey, R. D., & Palfrey, T. R. (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior, 10*, 6–38.

McKelvey, R. D., & Palfrey, T. R. (1998). Quantal response equilibria for extensive form games. *Experimental Economics, 1*, 9–41.

Mundt, P., Alfarano, S., & Milaković, M. (2015). Gibrat's law redux: Think profitability insterad of growth. *Industrial and Corporate Change, 25*(4), 549–571.

Mundt, P., & Oh, I. (2019). Asymmetric competition, rish, and return distribution. *Economics Letters, 179*, 29–23.

Muth, J. F. (1961). Rational expectations and the theory of price movements. *Econometrica, 29*(315-335).

Ömer, Ö. (2018). Dynamics of the us housing market: A quantal response statistical equilibrium approach housing market: A quantal response statistical equilibrium approach. *Entropy, 20*(11).

Porter, T. M. (1986). *Rise of statistical thinking, 1820-1900*. Princeton, NJ: Princeton University Press.

Scharfenaker, E., & Foley, D. K. (2017). Quantal response statistical equilibrium in economic interactions: Theory and estimation. *Entropy, 19*(444).

Scharfenaker, E., & Semieniuk, G. (2016). A statistical equilibrium approach to the distribution of profit rates. *Metroeconomica, 68*(3), 465–499.

Simon, H. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics, 69*(1), 99–118.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics, 50*(3), 665–690.

Smith, A. (1776 [1982]). *The wealth of nations.* London: Penguin.

Stanley, M., Amaral, L., Buldyrev, S., Havlin, S., Leschhorn, H., Maass, P., . . . Stanley, H. (1996). Scaling behavior in the growth of companies. *Nature, 379*(6568), 804–806.

Subbotin, M. (1923). On the law of frequency of error. *Mathematicheskii Sbornik, 31*, 269–301.

Suppes, P., Krantz, D. M., Luce, R. D., & Tversky, A. (1989). *Foundations of measurement: Geometrical, threshold, and probabilistic representations.* Academic Press Inc.

Vihola, M. (2012). Robust adaptive metropolis algorithm with coerced acceptance rate. *Statistics and Computing, 22*(5), 997–1008.

Wiener, N. (2019). *Diversity in segmention. patterns of immigrant competition in us labor markets* (Working Paper No. 01/2019). Department of Economics, The New School for Social Research.

Williams, M. A., Baek, G., Li, Y., Park, L. Y., & Zhao, W. (2017). Global evidence on the distribution of gdp growth rates. *Physica A, 468*, 750–758.

Wolpert, D. (2006). Information theory: The bridge connecting bounded rational game theory and statistical physics. In D. Braha, A. A. Minai, & Y. Bar-Yam (Eds.), *Complex engineered systems* (Chap. 12). Springer.

Yang, J. (2018). A quantal response statistical equilibrium model of induced technical change in an interactive factor market: Firm-level evidence in the eu economies. *Entropy, 20*(3).

8 Appendix A: Data

Data was gathered

9 Appendix B: QRSE Transformations

There are at least two useful coordinate transformations of 27 that make the QRSE model applicable to a wider range of possible economic problems. The first is to transform the kernel of 27 to coordinates based on the odds against the action $q[x] = e^{-\frac{x-\mu}{T}}$:

$$\hat{f}[q] \propto e^{\mathcal{H}[q]} e^{(\bar{\zeta} - \bar{T} \log[q]) \tanh\left[\frac{\log[q]}{2}\right]} \tag{30}$$

The second is to transform the kernel of 27 to the frequency of the action $p[x] = \frac{1}{1+e^{-\frac{x-\mu}{T}}}$:

$$\hat{f}[p] \propto e^{\mathcal{H}[p]} e^{\left(\bar{\zeta} - \bar{T} \log\left[\frac{1}{p}-1\right]\right) \tanh\left[\frac{1}{2} \log\left[\frac{1}{p}-1\right]\right]} \tag{31}$$

where $\bar{\zeta} = \gamma(\mu - \xi)$, and $\bar{T} = \gamma T$ become the only effective parameters in the model and

$$
\begin{aligned}
\mathcal{H}[p] &= -(1-p)\log[(1-p)] - p\log[p] \\
\mathcal{H}[q] &= -\frac{1}{1+q}\log[\frac{1}{1+q}] - \frac{q}{1+q}\log[\frac{q}{1+q}]
\end{aligned}
\tag{32}
$$

These transformation restrict the domain of $q \in [0, \infty]$ and $p \in [0, 1]$. The resulting marginal frequencies $\hat{f}[q]$ and $\hat{f}[p]$ approximate the behavior of the Generalized Beta (GB2) distribution and the Beta distribution respectively.
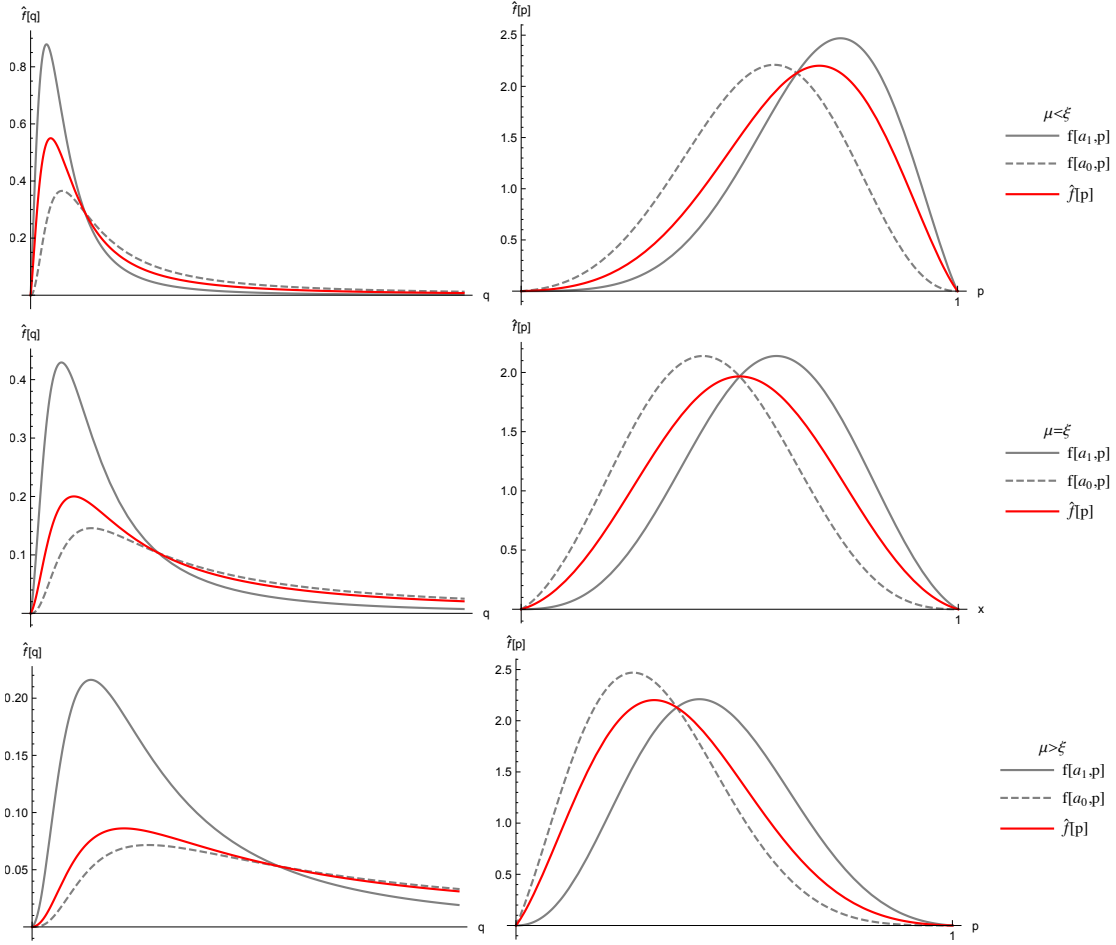
Figure 9: Marginal and joint frequency distributions transformed QRSE. Left: Log-odds of action, Right: Frequency of action. Top: $\bar{\zeta} = -2, \bar{T} = 1$, Middle: $\bar{\zeta} = 0, \bar{T} = 1$, Bottom: $\bar{\zeta} = 2, \bar{T} = 1$.