

Article

A Maximum Entropy Model of Bounded Rational Decision-Making with Prior Beliefs and Market Feedback

Benjamin Patrick Evans ^{*} and Mikhail Prokopenko 

Centre for Complex Systems, The University of Sydney, Sydney, NSW 2006, Australia;
mikhail.prokopenko@sydney.edu.au

* Correspondence: benjamin.evans@sydney.edu.au

Abstract: Bounded rationality is an important consideration stemming from the fact that agents often have limits on their processing abilities, making the assumption of perfect rationality inapplicable to many real tasks. We propose an information-theoretic approach to the inference of agent decisions under Smithian competition. The model explicitly captures the boundedness of agents (limited in their information-processing capacity) as the cost of information acquisition for expanding their prior beliefs. The expansion is measured as the Kullback–Leibler divergence between posterior decisions and prior beliefs. When information acquisition is free, the homo economicus agent is recovered, while in cases when information acquisition becomes costly, agents instead revert to their prior beliefs. The maximum entropy principle is used to infer least biased decisions based upon the notion of Smithian competition formalised within the Quantal Response Statistical Equilibrium framework. The incorporation of prior beliefs into such a framework allowed us to systematically explore the effects of prior beliefs on decision-making in the presence of market feedback, as well as importantly adding a temporal interpretation to the framework. We verified the proposed model using Australian housing market data, showing how the incorporation of prior knowledge alters the resulting agent decisions. Specifically, it allowed for the separation of past beliefs and utility maximisation behaviour of the agent as well as the analysis into the evolution of agent beliefs.

Keywords: decision-making; bounded rationality; complexity economics; information-theory; maximum entropy principle; quantal response statistical equilibrium

JEL Classification: D91; G41; D83; C61; C60; C50



Citation: Evans, B.P.; Prokopenko, M. A Maximum Entropy Model of Bounded Rational Decision-Making with Prior Beliefs and Market Feedback. *Entropy* **2021**, *23*, 669. <https://doi.org/10.3390/e23060669>

Academic Editors: Ryszard Kutner, Christophe Schinckus and Eugene Stanley

Received: 21 April 2021

Accepted: 21 May 2021

Published: 26 May 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Economic agents are often faced with partial information and make decisions under pressure, yet many canonical economic models assume perfect information and perfect rationality. To address these challenges, Simon [1] introduced bounded rationality as an alternate attribute of decision-making. Bounded rationality aims to represent partial access to information, with possible acquisition costs, and limited computational cognitive processing abilities of the decision-making agents.

Information theory offers several natural advantages in capturing bounded rationality, interpreting the economic information as the source data to be delivered to the agent (receiver) through a noisy communication channel (where the level of noise is related to the “boundedness” of the agent). This representation has spurred the creation of information-theoretic approaches to economics, such as Rational Inattention (R.I.) [2], and more recently, the application of R.I. to discrete choice [3]. Another approach represents decision-making as a thermodynamic process over state changes and employs the energy-minimisation principle to derive suitable decisions [4].

These approaches have shown how one can incorporate a priori knowledge into decision-making, but place no consideration to inferring these decisions based on observed

macroeconomic outcomes (e.g., a distribution of profit rates within a financial market) and market feedback loops. Independently, another recent information-theoretic framework, Quantal Response Statistical Equilibrium (QRSE) [5], was developed aiming to infer least biased (i.e., “maximally noncommittal with regard to missing information” [6]) decisions through the maximum entropy principle, given only the macroeconomic outcomes (e.g., when the choice data is unobserved). However, the ways to incorporate prior knowledge into such a system remain mostly unexplored.

In this work, we provide a unification of these approaches, showing how to incorporate prior beliefs into QRSE in a generic way. In doing so, we provide a least biased inference of decision-making, given an agent’s prior belief. Specifically, we show how the incorporation of prior beliefs affects the agent’s resulting decisions when their individual choices are unobserved (as is common in many real-world economic settings). The proposed information-theoretic approach achieves this by considering a cost of information acquisition (measured as the Kullback-Leibler divergence), where this cost controls deviations from an agent’s prior knowledge on a discrete choice set. When the cost of information acquisition is prohibitively high (i.e., when an agent is faced with limitations through time, cognition, cost, or other constraints), the agent falls back to their prior beliefs. When information acquisition is free, the agent becomes a perfect utility maximiser. The cost of information acquisition therefore measures the boundedness of the agent’s decision-making.

The proposed approach is general, allowing the incorporation of any form of prior belief, while separating the agents’ current expectations from their built-up beliefs. In particular, we show how incorporating prior beliefs into the QRSE framework allows for modelling decisions in a rolling way, when previous decisions “roll” into becoming the latest beliefs. Furthermore, we place the original QRSE in the context of related formalisms, and show that it is a special case of the general model proposed in our study, when the prior preferences (beliefs) are assumed to be uniform across the agent choices. Finally, we verify and demonstrate our approach using actual Australian housing market data, in terms of agent buying and selling decisions.

The remainder of the paper is organised as follows. Section 2 provides a background of information-theoretic approaches to economic decision-making, Section 3 describes QRSE and relevant decision-making literature. Section 4 outlines the proposed model, and Section 5 applies the developed model to the Australian housing market. Section 6 presents conclusions.

2. Background and Motivation

The use of statistical equilibrium (and more generally, information-theoretic) models remains a relatively new concept in economics [7]. For example, Yakovenko [8] outlines the use of statistical mechanics in economics. Scharfenaker and Semieniuk [9] detail the applicability of maximum entropy for economic inference, Scharfenaker and Yang [10] give an overview of maximum entropy and statistical mechanics in economics outlining the benefits of utilising the maximum entropy principle for rational inference, and Wolpert et al. [11] outline the use of maximum entropy for deriving equilibria with bounded rational players in game theory. Earlier, Dragulescu and Yakovenko [12] showed how in a closed economic system, the probability distribution of money should follow the Boltzmann-Gibbs law [13]. Foley [14] discusses Rational expectations and boundedly rational behaviour in economics. Harré [15] gives an overview of information-theoretic decision-theory and applications in economics, and Foley [16] analyses information-theory and results on economic behaviour.

Ömer [17] provides a comparison of “conventional” economic models and newly proposed ideas from complex systems such as maximum entropy methods and Agent-based models (ABM), which deviate from the assumption of homo economicus—a perfectly rational representative agent. Yang and Carro [18] discuss how a combination of agent-based modelling and maximum entropy models can be complementary, leveraging the analytical rigour of maximum entropy methods and the relative richness of agent-based modelling.

One of the key developments in this area is Quantal Response Statistical Equilibrium (QRSE) proposed by Scharfenaker and Foley [5]. This approach enabled applications of the maximum entropy method [6,19,20] to a broad class of economic decision-making. The QRSE model was further explored in [21], arguing that “any system constrained by negative feedbacks and boundedly rational individuals will tend to generate outcomes of the QRSE form”. The QRSE approach is detailed in Section 3.1.

Ömer [22–24] applies QRSE to housing markets (which we also use as a validating example), modelling the change in the U.S. house price indices over several distinct periods, and explaining dynamics of growth and dips. Yang [25] applies QRSE to a technological change, modelling the adoption of new technology for various countries over multiple years and successfully recovering the macroeconomic distribution of rates of cost reduction. Wiener [26–28] applies QRSE to labour markets, modelling the competition between groups of workers (such as native and foreign-born workers in the U.S.), and capturing the distribution of weekly wages. Blackwell [29] provides a simplified QRSE for understanding the behavioural foundations. Blackwell further extends this in [30], introducing an alternate explanation for skew, which arises due to the agents having different buy (enter) and sell (exit) preferences. Scharfenaker [31] introduces Log-QRSE for income distribution, and importantly, (briefly) mentions informational costs as a possible cause for asymmetries in QRSE. This is captured by measuring utility U as a sum $U[a, x] + C(a|x)$, allowing for higher costs (C) of entrance or exit into a market, where a is an action and x is a rate. Such a separation allows for an “alternative interpretation of unfulfilled expectations”.

These developments show the usefulness of maximum entropy methods, where we have placed particular focus on QRSE, for inferring decisions from only macro-level economic data. However, these approaches do not consider the contribution of a priori knowledge to the resulting decision-making process. The key objective of our study is to generalise the QRSE framework by the introduction of the prior beliefs, as well as the information acquisition costs as a measure of deviation from such priors.

3. Underlying Concepts

Two main concepts form the basis for the proposed model. The first is the QRSE approach developed by [5], and the second is a thermodynamics-based concept of decision-making derived from minimising negative free energy, proposed by [4].

3.1. QRSE

The QRSE framework aims to explain macroeconomic regularities as arising from social interactions between agents. There are two key assumptions stemming from the idea of Smithian competition: Agents observe and respond to macroeconomic outcomes, and agent actions affect the macroeconomic outcome, i.e., a feedback loop is assumed. It is this feedback that is deemed to cause the macroeconomic outcome to have a distribution that stabilises around an average value. Given only the macroeconomic outcome, QRSE infers the least biased distribution of decisions, which result in the observed macroeconomic distribution using the principle of maximum entropy. This makes QRSE particularly useful for inferring decisions when the individual decision level data is unobserved. In the following section, we outline the key notions behind QRSE [5].

3.1.1. Deriving Decisions

Agents are assumed to respond (i.e., make decisions) based on the macroeconomic outcome, for example, based on profit rates x . This is captured by the agents’ utility U . However, agents are assumed to act in a boundedly rational way, such that they may not always choose the option with the highest U , for example, if it becomes impractical to consider all outcomes. That is, agents are attempting to maximise their expected utility, subject to an entropy constraint capturing the uncertainty:

$$\max \sum_{a \in A} f[a|x] U[a, x] \quad (1)$$

$$\begin{aligned} &\text{subject to } \sum_{a \in A} f[a|x] = 1 \\ & - \sum_{a \in A} f[a|x] \log f[a|x] \geq H_{min} \end{aligned} \tag{2}$$

where $f[a|x]$ represents the probability of an agent choosing action a if rate x is observed. The first constraint ensures the probabilities sum to 1, while the second is a constraint on the minimum entropy. The minimum entropy constraint implies a level of boundedness such that there is some limit to the agents' processing abilities, which allows QRSE to deviate from perfect rationality.

Lagrange multipliers can be used to turn the constrained optimization problem of Equation (2) into an unconstrained one, which forms the following Lagrangian function:

$$\mathcal{L} = - \sum_{a \in A} f[a|x] U[a, x] - \lambda \left(\sum_{a \in A} f[a|x] - 1 \right) + T \left(- \sum_{a \in A} f[a|x] \log f[a|x] - H_{min} \right) \tag{3}$$

taking the first order conditions of Equation (3), and solving for $f[a|x]$ yields:

$$f[a|x] = \frac{1}{Z} e^{\frac{U[a,x]}{T}} \tag{4}$$

representing a choice of a mixed strategy by maximising the expected utility subject to an entropy constraint. This problem is dual to maximising entropy of the mixed strategy, subject to a constraint on the expected utility as detailed in Appendix A.1.

3.1.2. Deriving Statistical Equilibrium

From Section 3.1.1 we have a derivation for a decision function, where agents maximise expected utility subject to an entropy constraint introducing bounds in the agents processing abilities. In order to infer the statistical equilibrium based on observed macroeconomic outcomes, the joint probability $f[a, x]$ must be computed.

The joint distribution captures the resulting statistical equilibrium which arises from the individual agent decisions. While there are many potential joint distributions, using the principle of maximum entropy allows for inference of the least biased distribution. From an observer perspective, maximising the entropy of the model accounts for model uncertainty, by providing the maximally noncommittal joint distribution. To compute this, Scharfenaker and Foley [5] maximise the joint entropy with respect to the marginal probabilities (since individual action data is not available), by decomposing the joint entropy into a sum of the marginal entropy and the (average) conditional entropy.

The solution for $f[a|x]$, given by Equation (4), can be used to compute the joint probability $f[a, x]$, as long as marginal $f[x]$ is determined (since $f[a, x] = f[a|x]f[x]$). In order to derive $f[x]$, the approach considers the state dependant conditional entropy, represented as

$$H[A|x] = - \sum_{a \in A} f[a|x] \log f[a|x] \tag{5}$$

Scharfenaker and Foley [5] then use the principle of maximum entropy to find the distribution of $f[x]$ which maximises

$$\max_{f[x] \geq 0} H = - \int_x f[x] \log f[x] dx + \int_x f[x] H[A|x] dx \tag{6}$$

$$\begin{aligned} &\text{subject to } \int_x f[x] dx = 1 \\ & \int_x f[x] x dx = \zeta \end{aligned} \tag{7}$$

The first constraint ensures the probabilities sum to 1, and the second constraint applies to the mean outcome (with ζ being the mean from the actual observed data $\bar{f}[x]$). Importantly, there is also an additional constraint which models Smithian competition [32] in the market. Smithian competition models the feedback structure for competitive markets, for example, entrance into a market tends to lower the profit rates, and exit tends to raise the profit rates. This is captured as the difference between the expected returns conditioned on entrance, and the expected returns conditioned on exiting. This competition constraint can be represented as

$$\text{subject to } \int_x f[x](f[a|x] - f[\bar{a}|x])x dx = \delta \quad (8)$$

The combination of the conditional probabilities of Equation (4), which stipulate that the agents enter and exit based on profit rates, and the competition constraint of Equation (8) models a negative feedback loop that results in a distribution of the profit rates around an average (ζ).

Again, using the method of Lagrange multipliers, the associated Lagrangian becomes

$$\mathcal{L} = - \int_x f[x] \log f[x] dx + \int_x f[x] H[A|x] dx - \lambda \left(\int_x f[x] dx - 1 \right) - \gamma \left(\int_x f[x] x dx - \zeta \right) - \rho \left(\int_x f[x](f[a|x] - f[\bar{a}|x])x dx - \delta \right) \quad (9)$$

where taking the first order conditions of Equation (9), and solving for $f[x]$ yields

$$f[x] = \frac{1}{Z_A} e^{H[A|x] - \gamma x - \rho x(f[a|x] - f[\bar{a}|x])} \quad (10)$$

where Z_A is the partition function $Z_A = \int_x e^{H[A|x] - \gamma x - \rho x(f[a|x] - f[\bar{a}|x])} dx$. Note that in Equation (9) we use ρ as the Lagrangian multiplier for the competition constraint. Parameter ρ is referred to as β in [5], we have avoided this notation to avoid confusion with the thermodynamic β (inverse temperature) discussed in later sections.

Equations (4) and (10) comprise a fully defined joint probability. Crucially, QRSE allows for modelling the resultant statistical equilibrium even when the individual actions are unobserved—by inferring these decisions based on the principle of maximum entropy.

3.1.3. Limitations of Logit Response

In Section 3.1.1 we have seen how the logit response function used for decision-making in QRSE is derived from entropy maximisation. Following the Boltzmann distribution well known in thermodynamics, this logit response has seen extensive use throughout the literature arising in a variety of domains. For example, the logit function is used as sigmoid or softmax in neural networks, logistic regression, and in many applications in economics and game theory [33,34]. However, one important development not yet discussed is the incorporation of prior knowledge into the formation of beliefs. Up until now, we have considered a choice to be the result of expected utility maximisation based on entropy constraints from which the logit models have arisen. However, from psychology [35], behavioural economics [36,37], and Bayesian methods [38,39] we know that the incorporation of a priori information is often an important factor in decision-making. Thus, we explore the incorporation of prior beliefs into agent decisions in more detail in the following section (and the remainder of the paper).

Furthermore, one criticism of the logit response arises from the independence of irrelevant alternatives (IIA) property of multinomial logit models (which would extend to the conditional function used in QRSE in a multi-action case), which states that the ratio between two choice probabilities should not change based on a third irrelevant alternative. Initially, this may seem desirable, however, this can become problematic for correlated outcomes (of which many real examples possess). This criticism has been proved correct in

several thought experiment studies, showing violations of the IIA assumption [40]. The classical example is the Red Bus/Blue Bus problem [41,42].

Consider a decision-maker who must choose between a car and a (blue) bus, $A = \{\text{car, blue bus}\}$. The agent is indifferent to taking the car or bus, i.e., $p(\text{car}) = p(\text{blue bus}) = 0.5$. However, suppose a third option is added, a red bus which is equivalent to the blue bus (in all but colour). The agent is indifferent to the colour of the bus, so when faced with $A_1 = \{\text{blue bus, red bus}\}$ the agent would choose $p(\text{red bus}) = p(\text{blue bus}) = 0.5$. Now suppose the agent is faced with a choice between $A_2 = \{\text{car, blue bus, red bus}\}$. As per the IIA property, the ratio $\frac{p(\text{blue bus})}{p(\text{car})}$ (from $A, \frac{0.5}{0.5}$) must remain constant. So adding in a third option, the probability of taking any a becomes $p(a) = \frac{1}{3}$ (for all a), maintaining $\frac{p(\text{blue bus})}{p(\text{car})} = 1$. However, this has reduced the odds of taking the car from 0.5 to 0.33 based on the addition of an irrelevant alternative (i.e., the red bus in which the agent does not care about colour of the bus). In reality, the probability for taking the car should have stayed fixed at $p(\text{car}) = 0.5$, and the probability of taking a bus reduced to 0.25 each. This reduction in the probability of $p(\text{car})$ does not make sense for a decision-maker who is indifferent to the colour of the bus and is the basis for the criticism. This may not be immediately relevant for current QRSE models (especially binary ones), but with potential future applications, for example, in portfolio allocation, this could become an important consideration. For example, if adding an additional stock to a portfolio which is similar to an existing stock, it may not be desirable to reduce the likelihood of selecting other (unrelated) stocks.

3.2. Thermodynamics of Decision-Making

A thermodynamically inspired model of decision-making which explicitly considers information costs, as well as the incorporation of prior knowledge, is proposed by [4]. The proposed approach can be seen as a generalisation of the logit function, where the typical logit function can be recovered as a special case, but in the more general case manages to avoid the IIA property.

Ortega and Braun [4] represent changing probabilistic states as isothermal transformations. Given some initial state $x \in X$ with initial energy potential $\phi_0[x]$, the probability of being in state x is $p[x] = \frac{e^{-\beta\phi_0[x]}}{\sum_{x' \in X} e^{-\beta\phi_0[x']}}$ (from the Boltzmann distribution). Updating state to $f[x]$ corresponds to adding new potential $\Delta\phi_0[x]$. The transformation requires physical work, given by the free-energy difference $\Delta F[f]$. The free energy difference between the initial and resulting state is then

$$\begin{aligned} \Delta F[f] &= F[f] - F[p] \\ &= \sum_{x \in X} f[x] \Delta\phi_0(x) + \frac{1}{\beta} \sum_{x \in X} f[x] \log\left(\frac{f[x]}{p[x]}\right) \end{aligned} \quad (11)$$

which allows the separation of the prior $p[x]$ and the new potential $\Delta\phi_0[x]$. In economic sense, representing the negative of the new potential as the utility gain, i.e., $U(x) = -\Delta\phi_0[x]$, allows for reasoning about utility maximisation subject to an informational constraint, given here as the Kullback-Leibler (KL) divergence from the prior distribution [4]. Golan [43] shows how the KL-divergence naturally arises as a generalisation of Shannon entropy (of Equation (2)) when considering prior information, and Hafner et al. [44] show how various objective functions can be seen as functionally equivalent to minimising a (joint) KL-divergence, even those not directly motivated by the free energy principle. Such analysis makes the KL-divergence a logical and fundamentally grounded measure of information acquisition costs, captured as the divergence from a prior distribution.

Ortega and Stocker [45] then apply this formulation to discrete choice by introducing a choice set A (space of actions), which leads to the following negative free energy difference, for a given observation x :

$$-\Delta F[f[a|x]] = \sum_{a \in A} f[a|x] U[a, x] - \frac{1}{\beta} \sum_{a \in A} f[a|x] \log \left(\frac{f[a|x]}{p[a]} \right) \quad (12)$$

where again a represents a choice (or action), and U the utility for the agent. The first term of Equation (12) is maximising the expected utility, and the second term is a regularisation on the cost of information acquisition. Again, in this representation, information cost is measured as the KL-divergence from the prior distribution.

Taking the first order conditions of Equation (12) and solving for $f[a|x]$ yields

$$f[a|x] = \frac{p[a] e^{\frac{U[x,a]}{T}}}{\sum_{a' \in A} p[a'] e^{\frac{U[a',x]}{T}}} \quad (13)$$

where we have moved from inverse temperature β to temperature T for notational convenience, i.e., $T = \frac{1}{\beta}$. The key formulation here is the separation of the prior probability p from the utility gain (or the new potential from the initial potential). T then arises as the Lagrange multiplier for the cost of information acquisition (as opposed to the entropy constraint of QRSE, described in Section 3.1). We emphasise this aspect in later sections.

Revisiting the IIA property, the incorporation of the prior probabilities in Equation (A7) can adjust the choices away from the logit equation, and thus managing to avoid IIA. However, if desired, the free energy model reverts to the typical logit function in the case of uniform priors, and so this property can be recovered. In economic literature, a similar model is given by Rational Inattention (R.I.) by [2]. The relationship between R.I. and the free energy approach of [4,45] is detailed in Appendix C.

4. Model

In this section, we propose an information-theoretic model of decision-making with prior beliefs in the presence of Smithian competition and market feedback. Given an agent's prior beliefs and an observed macroeconomic outcome (such as the distribution of returns), the model can infer the least biased decisions that would result in such returns. Importantly, the incorporation of prior beliefs allows for reasoning about the decision-making of the agent based upon both their prior beliefs and their utility maximisation behaviour.

We develop upon the maximum-entropy model of inference from [5], and the thermodynamic treatment of prior beliefs formalised by [4], as outlined in Section 3.

4.1. Maximum Entropy Component

The proposed approach can be seen as a generalisation of QRSE, allowing for the incorporation of heterogeneous prior beliefs based on the free-energy principle. The key element is the information acquisition cost, measured as the KL-divergence which arises from the free-energy principle and has been shown to provide a fundamentally grounded application of Bayesian inference [46]. In order to derive decisions $f[a|x]$ for an action or choice a (e.g., buy, hold or sell) given an observed return x (e.g., a return on investment), we maximise the expected utility U subject to a constraint on the acquisition of information measured as the maximal divergence d between the posterior decisions and prior beliefs

$p[a]$. As mentioned, d is measured as the KL-divergence, which is the generalised extension of the original (Shannon) entropy constraint [43] introduced in Equation (2):

$$\begin{aligned} & \max \sum_{a \in A} f[a|x] U[a, x] \\ & \text{subject to } \sum_{a \in A} f[a|x] \log \left(\frac{f[a|x]}{p[a]} \right) \leq d \\ & \sum_{a \in A} f[a|x] = 1 \end{aligned} \quad (14)$$

The Lagrangian for Equation (14) then becomes

$$\mathcal{L} = \sum_{a \in A} f[a|x] U[a, x] - \lambda \left(\sum_{a \in A} f[a|x] - 1 \right) - T \left(\sum_{a \in A} f[a|x] \log \left(\frac{f[a|x]}{p[a]} \right) - d \right) \quad (15)$$

There are two distinct modelling views on such a formulation [47–50]. The first assumes that specific constraints are known from the data, for example, a maximal divergence d may be specified based on actual observations of agent behaviour. The second view, instead, would consider the Lagrange multiplier T to be a free parameter of the model, with the constraint d representing an arbitrary maximum value: Thus, this approach would optimise T in finding the best fit. In this work, we take the second perspective since underlying decision data is unavailable, and a specific restriction on divergent information costs should not be enforced. In other words, T is considered to be a free model parameter corresponding to different information acquisition costs, mapping to different (unknown) cognitive and information-processing limits d .

Looking at the final term in Equation (15), in the case of homogeneous priors, $\log p[a]$ is a constant which drops out of the solution, which is equivalent to the optimisation problem of Equation (3), and thus, recovers the original QRSE model. In the general case, the dependence on $\log(p[a])$ means that T instead serves as the Lagrange multiplier for the cost of information acquisition. Taking the first order conditions of Equation (15) and solving for $f[a|x]$ (as shown in Appendix A.2) yields

$$f[a|x] = \frac{1}{Z_{A|x}} p[a] e^{\frac{U[a,x]}{T}} \quad (16)$$

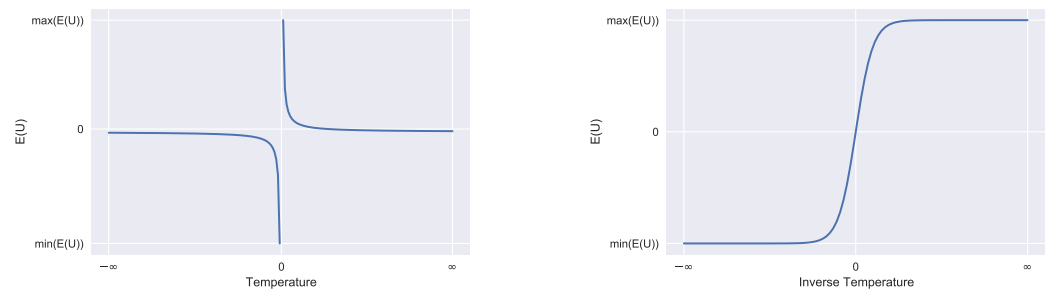
we see this as a generalisation of the logit function, which allows for the separation of the prior beliefs and the agent's utility function.

In the more general case, $p[a]$ can be heterogeneous for all a . Parameter T therefore controls the deviations from the prior (rather than from the base case of uniformity), that is, it controls the cost of information acquisition. Following [4], we observe the following limits

$$\begin{aligned} \lim_{T \rightarrow \infty} f[a|x] &= p[a] \\ \lim_{T \rightarrow 0, T \geq 0} f[a|x] &= e^{\frac{U[x,a]}{T}} = \max U[x, a] \\ \lim_{T \rightarrow 0, T < 0} f[a|x] &= e^{\frac{U[x,a]}{T}} = \min U[x, a] \end{aligned} \quad (17)$$

In the limit $T \rightarrow \infty$ (i.e., infinite information acquisition costs), the agent just falls back to their prior beliefs as it becomes impossible to obtain new information. In the limit $T \rightarrow 0$, the agent becomes a perfect utility maximiser (i.e., if information is free to obtain, the agent could obtain it all and choose the option that best maximises payoff with probability 1). In the $T < 0$ case, we see this corresponds to anti-rationality. For economic decision-making, we can limit temperatures to be non-negative, $T \geq 0$, although there are specific cases where such anti-rationality may be useful (e.g., modelling a pessimistic

observer or adversarial environments [4]). The relationship between temperature and utility is visualised in Figure 1.



(a) Temperate vs. Utility

(b) Inverse Temperate ($\frac{1}{T}$) vs. Utility

Figure 1. The effect of decision temperature T on the resulting expected payoffs (a). for the limits given by Equation (17). The inverse temperature $\frac{1}{T}$ (b) conveys the same information but may offer a more useful visualisation due to the continuity.

Crucially, large temperatures (costly acquisition) do not revert to the uniform distribution (as in the typical QRSE case, unless the prior is uniform), instead reverting to prior beliefs. This is visualised in Figure 2, and discussed in more detail in Section 4.3.

4.2. Feedback Between Observed Outcomes and Actions

Following [5], we use a joint distribution to model the interaction between the economic outcome x , and the action of agents a .

To recover a joint probability, we need to determine $f[x]$ (since $f[a, x] = f[a|x]f[x]$) which we do with the maximum entropy principle, as shown in Section 3.1. To do this, we maximise the joint entropy with respect to the marginal probabilities. That is,

$$\begin{aligned} \mathcal{L} = & - \int_x f[x] \log f[x] dx + \int_x f[x] H[A|x] dx - \lambda \left(\int_x f[x] dx - 1 \right) \\ & - \gamma \left(\int_x f[x] x dx - \xi \right) - \rho \left(\int_x f[x] \frac{p[a] e^{\frac{U[a,x]}{T}} - p[\bar{a}] e^{\frac{U[\bar{a},x]}{T}}}{Z_{A|x}} x dx - \delta \right) \end{aligned} \tag{18}$$

with

$$\begin{aligned} H[A|x] = & - \sum_{a \in A} f[a|x] \log f[a|x] \\ = & - \frac{1}{Z_{A|x}} \sum_{a \in A} p[a] e^{\frac{U[a,x]}{T}} \left(\log p[a] + \frac{U[a,x]}{T} - \log Z_{A|x} \right) \end{aligned} \tag{19}$$

An important point to be made here is that $H[A|x]$ still measures (Shannon) entropy. We have seen above how the new definition for $f[a|x]$ uses the KL-divergence as a generalised extension of entropy when incorporating prior information. In Equation (19), we do not use this divergence for an important reason. In Equation (14) we are measuring divergence from known prior beliefs, however, now when optimising Equation (18) we wish to infer decisions from unobserved decision data. This is where the principle of maximum entropy comes into play, i.e., we wish to maximise the entropy of our new choice data (which was derived from KL-divergence of prior beliefs), but we do not wish to perform cross-entropy minimisation as we do not have the true decisions $\tilde{f}[a|x]$. With this in mind, we still utilise the principle of maximum entropy as is done in QRSE for inference to obtain the least biased resulting decisions. This keeps the proposed extensions in the realm of QRSE, but comparisons to the principle of minimum cross-entropy [51,52] could be considered in future work particularly when some target distributions are known directly.

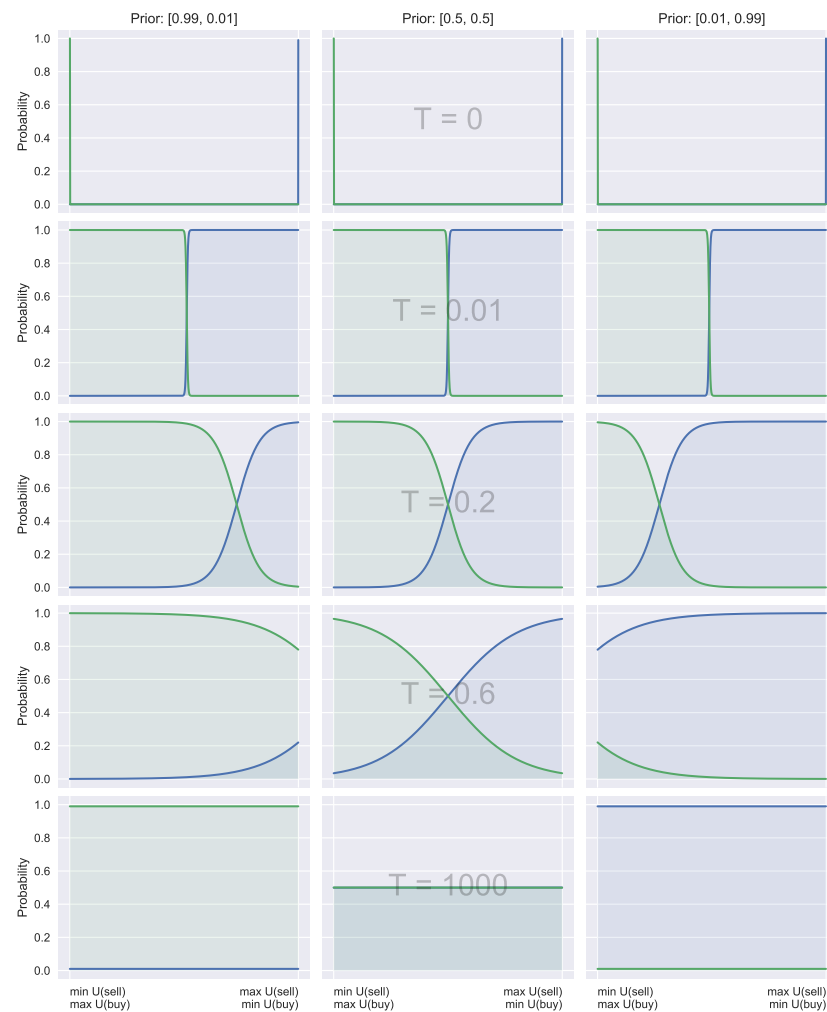


Figure 2. Decision Functions. All cases have equivalent utility functions. Each row has equivalent temperatures, showing how with matched parameters and utility, having an alternate prior can shift the decision-makers preference. Each column has different priors, given along the top of the first row to show how decision-makers decisions change based on their prior beliefs. On the left-hand side, preference is shifted towards the buying case. Likewise, on the right-hand side, preference is given to the selling case. The uniform case with equal preference is shown in the middle.

In Equation (18), ζ is known from the mean of the observed macroeconomic outcome, and so this constraint is used explicitly. This is in contrast to d (and δ) which are unknown as outlined in Section 4.1. The important distinction with Equation (18) is that the $f[a|x]$ functions (and $H[A|x]$) now use the updated expressions for $f[a|x]$, which incorporate the prior beliefs. Taking the partial derivative of \mathcal{L} with respect to $f[x]$, and solving for $f[x]$ gives

$$f[x] = \frac{1}{Z_A} e^{H[A|x] - \gamma x - \rho x \left(\frac{p[a]e^{-\frac{U[a,x]}{T}} - p[\bar{a}]e^{-\frac{U[\bar{a},x]}{T}}}{Z_{A|x}} \right)} \tag{20}$$

Equation (20) expresses the information acquisition cost in the form of the Lagrange multiplier T (from Equation (15)), and a competition cost in the form of the multiplier ρ .

As we have a solution for $f[a|x]$ (Equation (16)) and $f[x]$ (Equation (20)) in terms of prior beliefs and information acquisition costs, we can then derive all other probability

functions using the Bayes rule. That is, we can obtain $f[a, x]$, $f[x|a]$ and $f[a]$ which in turn incorporate these prior beliefs/acquisition costs:

$$\begin{aligned}
 f[a, x] &= f[a|x]f[x] \\
 &= \frac{p[a]e^{\frac{U[a,x]}{T} + H[A|x] - \gamma x - \rho x} \left(\frac{p[a]e^{\frac{U[a,x]}{T}} - p[\bar{a}]e^{\frac{U[\bar{a},x]}{T}}}{Z_{A|x}} \right)}{Z_{A|x}Z_A}
 \end{aligned}
 \tag{21}$$

We can obtain $f[a]$ by marginalising out x from the joint distribution:

$$\begin{aligned}
 f[a] &= \int_x f[a, x] \\
 &= \frac{1}{Z_A} \int_x \frac{1}{Z_{A|x}} p[a] e^{\frac{U[a,x]}{T} + H[A|x] - \gamma x - \rho x} \left(\frac{p[a]e^{\frac{U[a,x]}{T}} - p[\bar{a}]e^{\frac{U[\bar{a},x]}{T}}}{Z_{A|x}} \right)
 \end{aligned}
 \tag{22}$$

Finally, $f[x|a]$ can then be computed by a direct application of the Bayes rule: $f[x|a] = f[a, x]/f[a]$.

Given only an expected average value ζ (and the usual normalisation constraints), we have derived a joint probability distribution, which maximises the entropy subject to some information acquisition cost d , along with a competition cost δ . The resulting distribution free parameters (the Lagrange multipliers) are those which fit most closely to the true underlying distribution of returns. Thus, we have provided a generalisation of QRSE, which is fully compatible with the incorporation of prior beliefs.

4.3. Priors and Decisions

The introduced priors affect the conditional probabilities of agent decisions by shifting focus towards these preferred choices. The introduced priors allow the decision-maker to place more focus on particular actions if they have been deemed important a priori.

In Section 3.2 we showed how to separate the initial energy potential and new energy potential for distinguishing prior beliefs and utility functions. It is instructive to interpret these again as potentials, by setting $\alpha_a = T \log p[a]$, which allows us to represent the choice probability as

$$f[a|x] = \frac{1}{Z_{A|x}} e^{\frac{U[x,a] + \alpha_a}{T}}.
 \tag{23}$$

Equation (23) shows how α shifts the likelihood based on the prior preferences. An example of these shifts is visualised in Figure 2. This can be interpreted as placing more emphasis on actions deemed useful a priori as T increases. The information acquisition cost component T then controls the sensitivity between the utility and a priori knowledge, with a high T meaning higher dependence on prior information, and low T indicating a stronger focus on the utility alone.

The majority of binary QRSE models use a simple linear payoff definition for utility:

$$U[x, a] = x - \mu, \quad U[x, \bar{a}] = -(x - \mu).$$

With this definition, a tunable shift parameter μ serves as the expected fundamental rate of return. The relationship between μ and the real markets returns ζ (which was used as a constraint in Equation (7)), serves then as a measure of fulfilled expectations (i.e., if $\mu = \zeta$) or unfulfilled expectations ($\mu \neq \zeta$). This implies a symmetric shift parameter μ . As a specific example, if $a = \text{sell}$ and $\bar{a} = \text{buy}$, $\mu = 0.25$ means that at $x = 0.25$, buyers and sellers will be equally likely to participate in the market, i.e., $f[\text{sell}|\mu] = f[\text{buy}|\mu] = 0.5$. In this sense, μ can be seen as the indifference point. The symmetry arises from the fact that $f[\text{buy}|x] + f[\text{sell}|x] = 1$. Therefore, in the binary action case, it is possible to find a μ^* with the uniform priors $p = [0.5, 0.5]$ such that the decision functions will be equivalent

to μ with any arbitrary priors $p = [c, 1 - c]$, with $c \in [0, 1]$. In this sense, μ can be seen as encapsulating a prior belief.

However, explicit incorporation of prior beliefs on actions is useful here as it helps to separate the agents' expectations in relation to their prior belief (e.g., a higher μ resulted from needing to change from their past behaviour) and choose the actions for which an agent should emphasise acquiring more information. The introduced prior beliefs are strictly known before any inference is performed, whereas μ is the result of the inference process. The separation of prior beliefs and current expectations is important, as with μ alone this can not capture an agent's predisposition prior to performing any information processing. In addition, this applies more generally to any arbitrary utility functions (as QRSE is, of course, not limited to the linear shift utility function with μ outlined above), or when any preference is known about decisions a priori.

Consider also the three action case, $A = \{\text{buy, hold, sell}\}$, with the same utility functions as above but with the extra utility for holding being $U[x, \text{hold}] = 0$. We can see that it would be desirable if buying and selling no longer required this symmetry. The use of priors can introduce this asymmetry, by providing separate indifference points for buy/hold and sell/hold. Such asymmetry alters the resulting frequency distribution of transactions, and may help to explain various trading patterns [16]. The difference of symmetric and asymmetric buy and sell curves is shown in Figure 3. Figure 3 shows that such functions could be recovered by introducing a secondary shift parameter μ_2 . Parameter μ_1 (the original μ) then becomes the indifference point for buy and hold, and μ_2 for sell and hold. This is the method proposed in [30]. Introducing priors into this case again allows for separation of expectation μ , from prior belief and follows the same methodology as outlined above for the binary case. Furthermore, if we set $p[\text{hold}] = 0$, we recover the binary case. This highlights that the standard QRSE with binary actions and uniform priors is a special case of the ternary action case with heterogeneous priors.

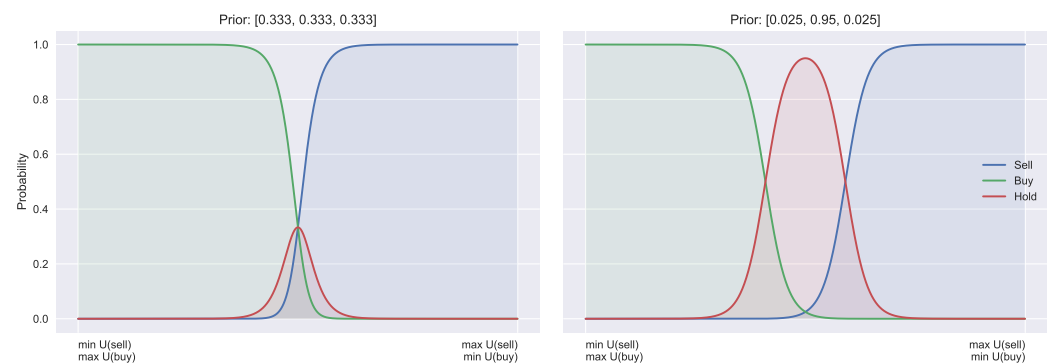


Figure 3. In the three-action case, the priors can introduce asymmetries by biasing the decision functions. This allows for separate indifferent points (right) vs. the uniform priors implying a single intersect (left).

From this, we can see how introducing priors alters the decision functions by allowing agents to focus on suitable a priori candidate actions. We have also shown how the binary case of a utility function with a shift parameter can be formalised to achieve equivalent results with a uniform prior and altered shift parameter. However, in the multi-action case, the priors allow for asymmetry, and in general, the priors may help with the optimisation process (by providing an alternate initial configuration). This approach also allows for the explicit separation of the two factors affecting an agent's choice, by distinguishing the contributions of prior beliefs and the utility maximisation.

4.4. Rolling Prior Beliefs

The proposed extension is general and allows for the incorporation of any form of prior beliefs, and in this section, we illustrate an example where the priors at time t are set as the resulting marginal probabilities from the previous time $t - 1$:

$$p_t[a] = f_{t-1}[a]$$

i.e., the prior belief $p_t[a]$ is set as the previous marginal probability $f_{t-1}[a]$ for taking action a (at $t = 0$, we use a uniform prior). Using the previous marginal probability as a prior introduces an “information-switching” cost, where T relates to the divergence from the previous actions, resulting in the following decision function:

$$f_t[a|x] = \frac{1}{Z_{A|x}} f_{t-1}[a] e^{\frac{U[a,x]}{T}}$$

That is, acquiring information on top of the previous knowledge comes at a cost (controlled by T). When the cost of information acquisition is high (large T), the agent reverts to the previously learnt knowledge (i.e., the marginal probabilities from $t - 1$). In contrast, when T is extremely small, the agent is able to acquire new information allowing deviation from their prior knowledge at $t - 1$. In the special case of $T = 0$, information is free, and the agent can become a perfect utility maximiser.

Given the expression for $f_t[a|x]$, we obtain the following solution for $f_t[x]$:

$$f_t[x] = \frac{1}{Z_A} e^{H[A|x] - \gamma x - \rho x} \left(\frac{f_{t-1}[a] e^{\frac{U[a,x]}{T}} - f_{t-1}[\bar{a}] e^{\frac{U[\bar{a},x]}{T}}}{Z_{A|x}} \right)$$

from which we can derive the joint and other probabilities, as shown in Section 4.1. This is exemplified in Section 5, in which we examine various priors for time-dependent applications.

5. Australian Housing Market

To exemplify the model, we use the Greater Sydney house price dataset provided by SIRCA-CoreLogic and utilised in [53,54]. This dataset is outlined in Appendix B. In [54], an agent-based model is used to explain and forecast house price trends and movement patterns as arising from the individual agent’s buy and sell decisions. Furthermore, the ABM implemented bounded rational agents driven by social influences (e.g., fear of missing out) and partial information about submarkets. While the resulting dynamics produced by the ABM accurately match the actual price trends, the decision-making mechanism and the bounded rationality of the agents were not theoretically grounded. In the following section, we aim to explain how the bounded rational behaviour of the agents operating in the housing market can be aligned with the model proposed in this study based on prior beliefs of agents and Smithian competition within the market. With this example, Smithian competition can be seen as agent decisions (buying or selling) affecting returns for an area, and agents decisions also being made based on returns for particular areas, i.e., a feedback loop is assumed in the market.

In particular, we want to explore what role an agent’s prior beliefs play in their resulting decisions. For example, given equivalent configurations (e.g., utility and returns) and different prior knowledge, how would the agent’s behaviour differ? Furthermore, we would like to explore the rationality of the agents, measured in terms of the cost of information acquisition, in order to see how the agents behave. For example, are agents predominantly reliant on past knowledge in times of market growth, resulting in unexpected downturns from mismanaged agent expectations? Alternatively, in deciding if it is a good time to buy or sell, the agents may balance their past knowledge with utility and current returns (i.e., the past knowledge would not be a predominant factor). The proposed model is particularly suited for answering such questions due to the low number

of free (and microeconomically) interpretable parameters, as well as the explicit separation of prior beliefs (as opposed to previous QRSE approaches). Our goal is not to infer the “best” prior, but rather to explore and compare dynamics resulting from various priors. In addition, we aim to verify the conjecture that during crises, and periods exhibiting non-linear market dynamics, macroeconomic conditions may become more heterogeneous, and thus, non-uniform priors may outperform uniform ones in such times.

5.1. Model

We use our model of binary actions with prior beliefs introduced in Section 4.1, with actions $A = \{\text{buy}, \text{sell}\}$. The decision functions are then given by

$$\begin{aligned} f_t[\text{buy}|x] &= \frac{1}{Z_{t,x}} p_t[\text{buy}] e^{\frac{U[x,\text{buy}]}{T}} \\ f_t[\text{sell}|x] &= \frac{1}{Z_{t,x}} p_t[\text{sell}] e^{\frac{U[x,\text{sell}]}{T}} \\ Z_{t,x} &= f_t[\text{buy}|x] + f_t[\text{sell}|x] \end{aligned} \quad (24)$$

where we explore a range of p_t (prior at time t) functions, discussing their effects on decision-making and resulting probability distributions.

5.1.1. Priors

While the proposed approach is capable of incorporating any form of prior belief on the choice set A , below we outline several example priors which we explore. In exploring these priors, we highlight differences in resulting agent posterior decisions based on various prior beliefs.

Uniform

We begin with a uniform prior. The uniform probability represents the default case of QRSE, where each action has an equally weighted prior. In the binary case, this corresponds to $p_t[a] = 0.5$ for all t and a . This corresponds to an agent who is agnostic to the available actions before observing U .

Previous

Next we look at a “previous” prior. The previous prior uses the marginal action probabilities from the previous time step as priors to the current timestep. This means at time t , $f_t[a]$ plays the role of a posterior probability of making a decision, however, at time $t + 1$ $f_t[a]$ now serves as the empirical prior. This is the example introduced in Section 4.4. This corresponds to $p_t[a] = f_{t-1}[a]$ for $t > 0$, and $p_t[a] = 0.5$ for $t = 0$. The previous prior represents an empirical prior where the decision is conditioned on previous market information, where T controls the level of influence from the previous market stage (in our case, each year). A high T means high influence from the past market state, whereas low T means focusing on current market conditions alone (as measured by U). In the extreme case of $T = \infty$, a backward looking expectations [55] approach is recovered where decisions are assumed to be a function purely of past decisions, however, in the more general case with $T < \infty$, U adjusts the decisions based on the current market state.

Mean

We also consider a mean prior. The mean prior uses the average marginal action probability from all previous timesteps. This corresponds to $p_t[a] = \frac{\sum_{t'=0}^{t-1} f_{t'}[a]}{t}$, for $t > 0$, and $p_t[a] = 0.5$ for $t = 0$. This can be seen as belief evolution, where over time, the previous decisions help build the current prior (modulated by T) at each stage.

Extreme Priors

As two further examples, we introduce extreme priors (more for visualisation/discussion sake as opposed to being particularly useful). The extreme buy prior corresponds to a strong prior preference for the buy action, $p_t[\text{buy}] = 0.99, p_t[\text{sell}] = 0.01$, for all t . Likewise, the extreme sell case is simply the inverse of the buy case, a strong prior preference for selling, i.e., $p_t[\text{sell}] = 0.99, p_t[\text{buy}] = 0.01$, for all t .

However, the formulations provided above by no means represent an exhaustive set of possible priors. For example, Genewein et al. [56] discuss “optimal” priors, which draws parallels with rate-distortion theory and can be seen as building abstractions of decisions (see Appendix C). Adaptive expectations [57] are discussed in [58–60], where priors could be partially adjusted based on some strength term (λ_E), where the strength term adjusts the contribution from some error. For example, an adaptive prior could be represented as $p_t = p_{t-1} + \lambda(p_{t-1} - \hat{p}_{t-1})$, where \hat{p}_{t-1} is the actual known likelihood of actions from the previous time period. With our specific housing market data, we do not have \hat{p} , i.e., we do not have the true buying and selling likelihoods, but if known, such information could be used to adjust future beliefs, i.e., over time the adaptive priors would adjust decisions based on the previously observed likelihoods (controlled by λ). The proposed approach makes no assumption about the forms of prior beliefs, so the ideas outlined above can be incorporated into the method outlined here by adjusting the definition of p_t .

5.2. Results

We fit the distributions with the various priors outlined in Section 5.1.1 to the actual underlying return data, to estimate how well we are able to capture this distribution and explore the effects that these priors have on the resulting distribution. The results are presented in Table 1, which summarises the likelihood and the percentage of the explained variability (measured as Information Distinguishability (I.D.) [61]) compared to the underlying distribution. We see that there are no large differences in general between the priors in terms of the explained variability. However, the goal here is not to argue for the “best” prior fitting the dataset in terms of the explained variability, but rather to explore differences in the agent behaviour based on the prior knowledge (using the housing dataset as an example). Thus, the resulting fitted distributions $f[x]$, which are visualised in Figure A5, are more interesting. We observe how altering prior beliefs result in different resulting distributions and discuss how the incorporation of prior beliefs allows for a separation of the agents’ utility maximisation behaviour from their previous knowledge. From Figure A5 we can also see how the priors can alter the optimisation process, for example, a good (bad) prior may help (harm) the optimisation by providing alternate initial configurations. The extreme priors can be seen as harmful, for example, in 2012 where the resulting distributions are unable to capture the true underlying distribution. The reason for this is being unable to find suitable T to enable appropriate divergence from the extreme prior beliefs. In contrast, well selected priors can help the optimisation process and result in better fitting distributions, such as in 2016 where the decisions resulting from the mean and previous prior fit the true data significantly better than the uniform prior.

The agents’ decision functions $f[a|x]$ are visualised in Figure A7 which makes it clear how each prior adjusts the resulting probability of taking an action (and thus, alters the decisions). From this, we can see different probabilistic behaviours despite having equivalent utility functions and optimisation processes due to varying prior beliefs. For example, with the extreme priors, we observe a clear shift towards the strongly preferred action.

Figure A6 shows the resulting joint distributions $f[a, x]$, combining the results of Figures A5 and A7, since $f[a, x] = f[a|x]f[x]$. Looking at the second row of each plot in Figure A6, we can see a visual representation of how the joint probabilities adjust over time when using the previous year as the prior belief.

Table 1. Resulting likelihood and percentage of variability explained for each year, when compared to the actual underlying distribution (i.e., those given in Figure A2). Optimisation is done by minimising the negative log-likelihood between the resulting distributions and the actual distribution of returns.

	Uniform	Previous	Mean	Extreme Buy	Extreme Sell
2006	1082 (93%)	1082 (93%)	1082 (93%)	885 (59%)	1005 (74%)
2007	1089 (92%)	1089 (92%)	1090 (90%)	939 (68%)	1042 (83%)
2008	998 (95%)	905 (78%)	998 (95%)	998 (95%)	998 (95%)
2009	918 (96%)	918 (96%)	866 (88%)	880 (85%)	875 (85%)
2010	857 (95%)	857 (95%)	857 (95%)	740 (62%)	857 (95%)
2011	1045 (92%)	1044 (91%)	1047 (92%)	1045 (91%)	873 (62%)
2012	1067 (96%)	1067 (96%)	1067 (96%)	162 (6%)	142 (8%)
2013	1080 (90%)	1076 (90%)	1083 (90%)	983 (77%)	1075 (91%)
2014	938 (98%)	851 (74%)	938 (98%)	875 (71%)	938 (98%)
2015	860 (96%)	860 (96%)	860 (96%)	33 (10%)	808 (71%)
2016	873 (84%)	932 (95%)	908 (86%)	817 (70%)	932 (95%)
2017	916 (97%)	916 (97%)	916 (97%)	812 (76%)	916 (97%)
2018	989 (88%)	932 (85%)	933 (85%)	955 (82%)	998 (91%)
2019	1101 (92%)	1103 (92%)	1067 (94%)	1101 (92%)	952 (76%)

The resulting marginal action probabilities are visualised in Figure 4, where we observe clear market peaks and dips which match the actual returns of Figure 5, aligning with the general trends observed in Figure A1. The priors work on either increasing or decreasing the resulting marginal probabilities. For example, in the extreme sell case we see much higher resulting probabilities for $f[\text{sell}]$, likewise in the extreme buying case, we see much higher probabilities for $f[\text{buy}]$. The general peaks/dips remain in both cases. Overall, this shows how the prior belief can influence the resulting marginal probabilities.

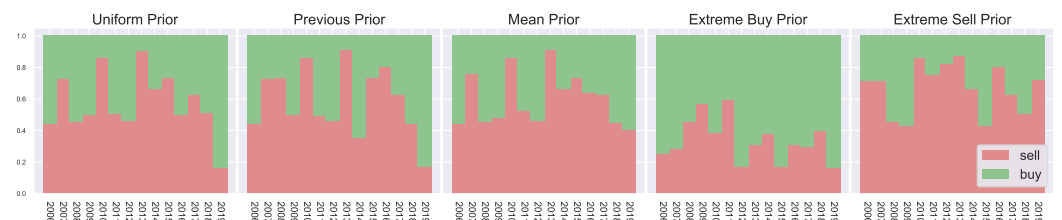


Figure 4. Resulting marginal probabilities $f[a]$ for varying priors. Green represents $f[\text{buy}]$, and red represents $f[\text{sell}]$.



Figure 5. Real Average Returns.

Using the previous year’s marginal probability as a prior for the current year has a smoothing effect on the resulting year-to-year marginal probabilities. Comparing the previous prior with the uniform prior in Figure 4, we observe, particularly during 2015–2018, a more defined/well-behaved step-off in $f[\text{sell}]$. This indicates the slowing of returns during these years. At the same time, the uniform priors are more affected by local noise, potentially overfitting to only the current time period, since no consideration can be given to

the past behaviour of the market. This results in larger fluctuations in the agent behaviour as they have no concept of market history.

5.3. Role of Parameters

One of the benefits of QRSE is the low number of free parameters which results in a relatively interpretable model. There are four free parameters in the typical QRSE distribution: T, μ, ρ and γ , each with a corresponding microeconomic foundation. In this section, we discuss the two main parameters of interest in this work: The decision temperature T and agent expectations μ , and the effect that prior beliefs have on the resulting values (and interpretation) of these parameters. We also include discussion on the impact of decisions on resulting outcomes ρ and skewness of the resulting distributions γ in Appendix D, since ρ and γ were less affected by the introduced extensions. There is an additional parameter ξ (shown in Figure 5), which is not a free parameter, representing the mean of the actual returns and serving as a constraint on the mean outcome in Equation (7).

5.3.1. Decision Temperature

The decision temperature T controls the level of rationality and deviations from an agent's prior beliefs. An extremely high temperature corresponds to high information acquisition cost and results in choosing actions simply based on the prior belief. In contrast, an extremely low temperature corresponds to utility maximisation, and in the case of free information ($T = 0$) a perfect utility maximiser is recovered (i.e., homo economicus). In the housing example used here, T relates to the ability of an agent to learn all the required knowledge of the market, i.e. the actual profit rates for various areas. With $T = 0$, the agent has perfect knowledge of the current market profitability. With $T > 0$, this represents some friction with acquiring such information, e.g., it can be difficult to gather all the required information to make an informed choice due to, for example, search costs. From a psychological perspective, T can be a measure of the "just-noticeable difference" [62], meaning microeconomically, T is related to the ability of an agent to observe quantitative differences in resulting choices. High T means the agent is unable to distinguish choices based on U , due to high information-processing costs, so instead acts according to their previously learnt knowledge.

Since T is related to the prior, we see differences in the resulting values visualised in Figure 6. What can be observed from looking at the general trends of T is that it peaks in the years with high average growth (large ξ), such as 2015, as these years correspond to a growing market, and agents require less attention to market conditions, although this depends on the prior used.



Figure 6. Decision Temperature.

Looking at the previous marginal probability as the prior (the orange profile), we observe in the build-up phase to 2015 increasing decision temperatures corresponding to agents acting on these previous beliefs. As these beliefs were also positive (i.e., agents expected favourable returns), these large returns can be explained by the agents continuously expecting this growth. This pattern changed in 2016, when the market "reverses": Now

the agents must focus instead on their current utility since their prior beliefs no longer reflect the current market state. Such market reversals are categorised by low decision temperatures, since using the previous action probabilities now becomes misinformative (in contrast to the “building”/trend-following stages). This indicates an increased focus on agent rationality in times of market reversals. The incorporation of prior beliefs (particularly using the previous priors) is useful as it allows for the discussion to be extended in the temporal sense (as is done here). In other words, we can consider “building” the agent’s beliefs as possible underlying causes for market collapses and relating the rationality of agents to the relative state of the market.

5.3.2. Agent Expectations

In microeconomic terms, parameter μ captures the agent’s expectations. A large μ corresponds to an optimistic agent, who is expecting high returns from the market. In contrast, a low μ corresponds to a pessimistic agent, who is expecting poor returns from the market. As this works to shift the decision functions, there is a relation between the prior and parameter μ , since the prior also works as shifting preferences towards a priori preferred actions as shown in Section 4.3. There is also a relationship between μ and γ (outlined in Appendix D.2), since γ can help to account for unfulfilled agent expectations by adjusting the skew of the resulting distributions.

Generally, the agent’s beliefs are within the $\pm 2.5\%$ range (expecting between a 2.5% quarterly growth or 2.5% dip), which corresponds to the bulk of the area under the curve in Figure A2. This means that the agent’s expectations develop in accordance with actual market conditions, as can be seen in Figure 7.

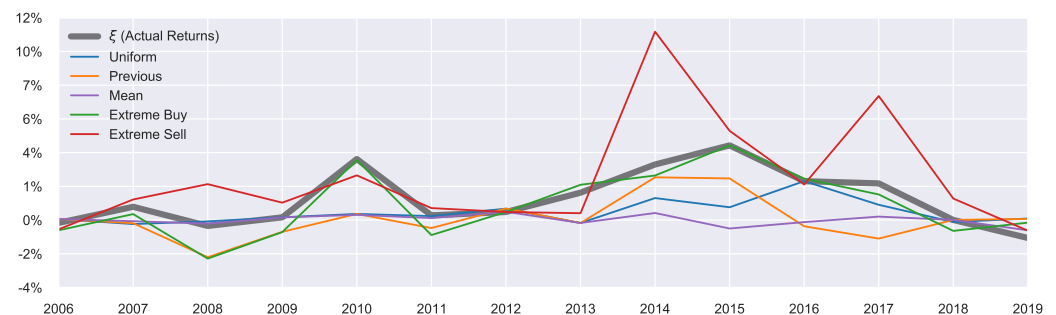


Figure 7. Agent Expectations vs. Actual Returns (in black).

The extreme priors result in larger absolute values of μ since larger shifts are needed to offset the (perhaps) poor prior beliefs. This can be seen in 2014 particularly, where the extreme sell prior has $\mu = 10\%$.

The values of previous prior μ tend to have a larger magnitude than the uniform priors, since as mentioned, these priors can capture build-up of beliefs (and as such some “trend-following” can be captured). For example, the year 2008 saw the lowest average returns ξ , as shown in Figure 5. Using the previous prior, the agents’ expectations correctly match the sign of the actual returns in 2008 (i.e., agents correctly expected a decline in house prices). This results in more pessimistic agents than those using the uniform prior since they can reflect on the market performance from 2007. Likewise, during 2013–2015, the values of previous prior μ become larger than those for the uniform prior, since they are building on the previous years expectations which were all positive. In contrast, the period 2015–2017 saw a steady decline in agents expectations of returns with previous priors, reflecting the overall market state which appeared to be in a downward trend. The previous priors were able to capture this trend. Using the uniform priors, the year 2016 had a higher μ than the market peak of 2015. The reason is that uniform priors are unable to capture the fact that the previous timestep had higher (or lower) returns than the current timestep. In this case, the discussion can not be extended in the temporal sense of “building” on beliefs, and agents may miss such crucial temporal information without the incorporation of prior

beliefs. This is evidenced by the significantly lower performance of the uniform prior in 2016 in comparison to the previous prior, as shown in Table 1, highlighting the usefulness of non-uniform (and temporal-based) priors in times of market crises and reversals.

5.4. Temporal Effects of Data Granularity on Decisions

In Section 5.2, we have analysed agent decisions over the previous 15 years, where decisions were grouped annually. This level of granularity was chosen to examine different agent behaviour from year to year. However, other levels of grouping can also be explored to give an insight into the impact of noise on the inference process. For example, an extremely granular grouping will likely result in additional noise in the decision-making process, which may or may not be impacted by the incorporation of prior beliefs. Likewise, a low granular grouping can be seen as “pre-smoothed”, which may work in a similar fashion to the incorporation of prior temporal-based beliefs at a higher granularity, which we have seen can smooth the resulting decisions. In this section, we examine the usefulness of prior beliefs in such situations, providing comparisons with alternate data representations.

Two additional levels of granularity are considered, one more granular and one less granular than the annual groupings introduced in Section 5.2. We look at quarterly data, as well as aggregate groupings based on market state. In doing so, we have three levels for categorising agent behaviour: Quarterly, annually, and aggregated market state. This allows us to compare resulting agent decisions across different temporal scales, comparing the differences generated by the incorporation of prior beliefs and various data-level modifications.

The aggregate market state data groups years into “terms”, which correspond with various “stages” of the market. These are growth and crash phases, highlighted as “Pre Crash” (Mid 2006–2007), “Crash” (2008), “Recovery 1” (2009–Mid 2011), “Small Crash” (Mid 2011–Mid 2012), “Recovery 2” (Mid 2012–Mid 2018) and “Recent Crash” (Mid 2018 to 2020). The overall market trends can be visualised in Figure A1 to see market returns for each corresponding “term”.

The resulting decision likelihoods $f[A]$ are presented in Figure 8. In analysing the differences in resulting marginal probabilities between the various granularities, we can observe the impact from data-level modifications, i.e., performing inference on a larger time scale for macroeconomic observations, and how the incorporation of prior information affects such results. In Section 5.2 we have mentioned the previous and mean priors can have a smoothing effect on resulting decisions, in this sense, the lower granularity groupings (the market state based grouping) can also be seen as a smoothed version of the macroeconomic outcomes, i.e. pre-smoothing the data by considering a much larger interval composed of several years for groupings. We see that the incorporation of prior information helps preserve some important information in such settings. Looking at the left-most column of Figure 8 (the uniform priors), we can see the overall “shape” of the peaks and dips in preferences $f[a]$ is lost with aggregate groupings. For example, in the quarterly breakdown, there is a clear preference for selling in the later region in the range 2014–2017, corresponding to the highest growing market, which is labelled as “Recovery 2” in the aggregated version. When considering the “Recovery 2” with uniform priors, such a clear preference is lost, and the “Pre Crash” and “Initial Recovery” have a higher corresponding preference. This is because the agents can not separate past market information from the current market state and act purely based on the current utility. In contrast, with both the mean and the previous prior, such overall trends are preserved across the various granularities since agents can distinguish favourable environments when compared with previous market states (as captured by their prior beliefs). This additional temporal insight provides an important consideration and shows that even with various data-level smoothing or preprocessing (i.e., considering alternate data groupings) the prior information remains useful and highlights various market states and corresponding agent preferences.



Figure 8. $f[a]$ for varying granularities.

A key takeaway from this exploration is that the potential for temporal analysis introduced by the prior beliefs provides additional insights into decision-making. These insights can not be generated by simple data-level modifications. Furthermore, the decision temperature T provides a way to modulate market state changes when considering agent decision-making.

6. Discussion and Conclusions

Despite many well-founded doubts of perfect rationality in decision-making, agents are often still modelled as perfect utility maximisers. In this paper, we proposed an approach for inference of agent choice based on prior beliefs and market feedback, in which agents may deviate from the assumption of perfect rationality.

The main contribution of this work is a theoretically grounded method for the incorporation of an agent's prior knowledge in the inference of agent decisions. This is achieved by extending a maximum entropy model of statistical equilibrium (specifically, Quantal Response Statistical Equilibrium, QRSE), and introducing bounds on the agent processing abilities, measured as the KL-divergence from their prior beliefs. The proposed model can be seen as a generalization of QRSE, where prior preferences across an action set do not necessarily have to be uniform. However, when uniform prior preferences are assumed, the typical QRSE model is recovered. The result is an approach that can successfully infer least biased agent choices, and produce a distribution of outcomes matching that of the actual observed macroeconomic outcomes when individual choice level data is unobserved.

In the proposed approach, the agent rationality can vary from acting purely on prior beliefs, to perfect utility maximisation behaviour, by altering the decision temperature. Low decision temperatures correspond to rational actors, while high decision temperatures represent a high cost of information acquisition and, thus, revert to prior beliefs. We showed how varying an agent's prior belief altered the resulting decisions and behaviour

of agents, even those with equivalent utility functions. Importantly, the incorporation of prior beliefs into the decision-making framework allowed the separation of two key elements: The agent's utility maximisation, and the contribution of the agent's past beliefs. This separation allowed for a discussion on the decision-making process in a temporal sense, being able to refer to the previous decisions. This allows for investigation into the building of beliefs over time, elucidating resulting microeconomic foundations in terms of the underlying parameters.

It is worth pointing out some parallels with, and differences from, the frameworks of embodied intelligence and information-driven (guided) self-organisation, in which embodiment is seen as a fundamental principle for the organisation of biological and cognitive systems [63–66]. Similar to these approaches, we consider information-processing as a dynamic phenomenon and treat information as a quantity that flows between the agent and its environment. As a result, an adaptive decision-making behaviour emerges from these interactions under some constraints. Maximisation of potential information flows is often proposed as a universal utility for such emergent agent behaviour, guiding and shaping relevant decisions and actions within the perception-action loops [67–70]. Importantly, these studies incorporate a trade-off between minimising generic and task-independent information-processing costs and maximising expected utility, following the tradition of information bottleneck [71].

In our approach, we instead consider specific information acquisition costs incurred when the agents need to update their relevant beliefs in the presence of (Smithian) competition and market feedback. The adopted thermodynamic treatment of decision-making allows us to interpret relevant economic parameters in physical terms, e.g., agent's decision temperature T , the strength of negative feedback ρ , and skewness of the resulting energy distribution γ . Interestingly, the decision temperature appears in our formalism as the Lagrange multiplier of the information cost incurred when switching posterior and prior beliefs (KL-divergence). The KL-divergence can be interpreted as the expected excess code-length that is needed if a non-optimal code that was optimal for the prior (outdated) belief is used instead of an optimal code based on the posterior (correct) belief. Thus, the decision temperature modulates the inference problem of determining the true distribution given new evidence, in a forward time direction [72]. Moreover, the thermodynamic time arrow (asymmetry) is maintained only when decision temperatures are non-zero.

We demonstrated the applicability of the method using actual Australian housing data, showing how the incorporation of prior knowledge can result in agents building on past beliefs. In particular, the agent focus can be shown to shift from utility maximisation to acting on previous knowledge. In other words, during the periods when the market has been performing well, the agents were shown to become overly optimistic based on the past performance.

The generality of the proposed approach makes it useful for incorporating any form of prior information on the agent's choice set. Moreover, we have shown that the default QRSE is a special case of the proposed extension with uniform (i.e., uninformative) priors. Therefore, the proposed approach can be seen as an extension of QRSE, which accounts for prior agent beliefs based on information acquisition costs. As the QRSE framework continues to be expanded, the generalised model proposed here could become an important approach. Particularly, this would be useful whenever prior knowledge on agent decisions is known, as well as in multi-action cases when the IIA property of the general logit function is undesirable. Other relevant applications include scenarios with multiple time periods, allowing for a detailed temporal analysis and exploration of the cost of switching between equilibria (measured as an information acquisition cost from prior beliefs).

Author Contributions: B.P.E. and M.P.; Funding acquisition, M.P.; Software, B.P.E.; Supervision, M.P.; Writing—original draft, B.P.E. and M.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Australian Research Council Discovery Project DP170102927.

Data Availability Statement: The real-estate pricing data used in this work were made available under license for this study by SIRCA-CoreLogic (<https://www.corelogic.com.au/industries/residential-real-estate>).

Acknowledgments: The authors would like to thank Kirill Glavatskiy and Michael S. Harré for many helpful discussions regarding the Australian housing market, as well as Adrián Carro, Jangho Yang and anonymous reviewers for various comments. The authors would also like to acknowledge the Securities Industry Research Centre of Asia-Pacific (SIRCA) and CoreLogic, Inc. (Sydney, Australia) for their data on Greater Sydney housing prices.

Conflicts of Interest: The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results.

Appendix A. Derivations

Appendix A.1. Decision Duality

There are two main perspectives, the first is of the agent performing actions within the system, and the second is of the system observer [29].

Each of the two perspectives allows to capture the uncertainty faced by either the actor or the observer, by imposing a constraint on entropy. In this section, we outline the duality that arises from these perspectives, showing that a duality exists between maximum entropy models, and entropy constrained models [7]. Additional discussion on such perspectives is given in [21].

Modelling the actor corresponds to maximising the expected utility subject to a fixed entropy constraint. This is the method outlined in Section 3.1.1. In this case, the agent can be seen as a boundedly rational decision-maker, in that they might not have all of the information required to make a perfectly rational choice.

The alternate perspective, modelling an observer, corresponds to maximising the entropy of the decisions subject to a fixed expected utility. With this perspective, we capture modelling uncertainty from the observer. The observers problem is formulated as follows

$$\begin{aligned} \max & - \sum_{a \in A} f[a|x] \log f[a|x] \\ \text{subject to} & \sum_{a \in A} f[a|x] = 1 \\ & \sum_{a \in A} f[a|x] U[a, x] \geq U_{min} \end{aligned} \quad (\text{A1})$$

where U_{min} represents the minimum expected utility. In order to see the duality of Equations (A1) and (1), we formulate the following Lagrangian for converting Equation (A1) into an unconstrained optimization problem.

$$\mathcal{L} = - \sum_{a \in A} f[a|x] \log f[a|x] - \lambda \left(\sum_{a \in A} f[a|x] - 1 \right) + \beta \left(\sum_{a \in A} f[a|x] U[a, x] - U_{min} \right) \quad (\text{A2})$$

where again, taking the first order conditions and solving for $f[a|x]$ yields

$$f[a|x] = \frac{1}{Z} e^{\beta U[a, x]} \quad (\text{A3})$$

We can see Equation (A3) is equivalent with Equation (4) with $\beta = \frac{1}{T}$, which highlights an important dualism between the two perspectives.

Appendix A.2. Decision Function

By setting the partial derivative of the unconstrained optimisation problem given in Equation (15) with respect to $f[a|x]$ to 0, we can obtain the following definition for $f[a|x]$:

$$\begin{aligned} \frac{d\mathcal{L}}{df[a|x]} &= U[a, x] - \lambda - T \log\left(\frac{f[a|x]}{p[a]}\right) = 0 \\ f[a|x] &= e^{\frac{U[a, x]}{T} - \lambda + \log p[a]} \end{aligned} \quad (\text{A4})$$

and, using the normalisation constraint $\sum_{a \in A} f[a|x] = 1$, we obtain the following decision function

$$\begin{aligned} f[a|x] &= \frac{1}{Z_{A|x}} e^{\frac{U[a, x]}{T} + \log p[a]} \\ &= \frac{1}{Z_{A|x}} p[a] e^{\frac{U[a, x]}{T}} \end{aligned} \quad (\text{A5})$$

with the partition function $Z_{A|x} = \sum_{a' \in A} p[a'] e^{\frac{U[a', x]}{T}}$.

Appendix B. Australian Housing Market Data

Data from 2006–2020 is used. Data is split into individual years. We use the rolling median price for each area and then measure the quarterly percentage growth rate for the areas. The month-to-month percentage changes are visualised in Figure A1. The distributions of the returns are visualised in Figure A2.

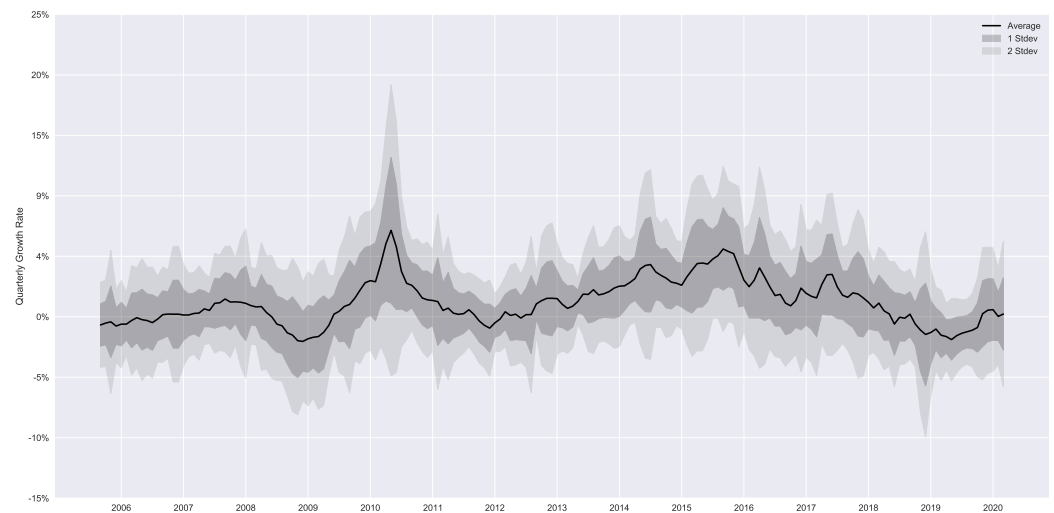


Figure A1. Quarterly returns in the Sydney housing market.

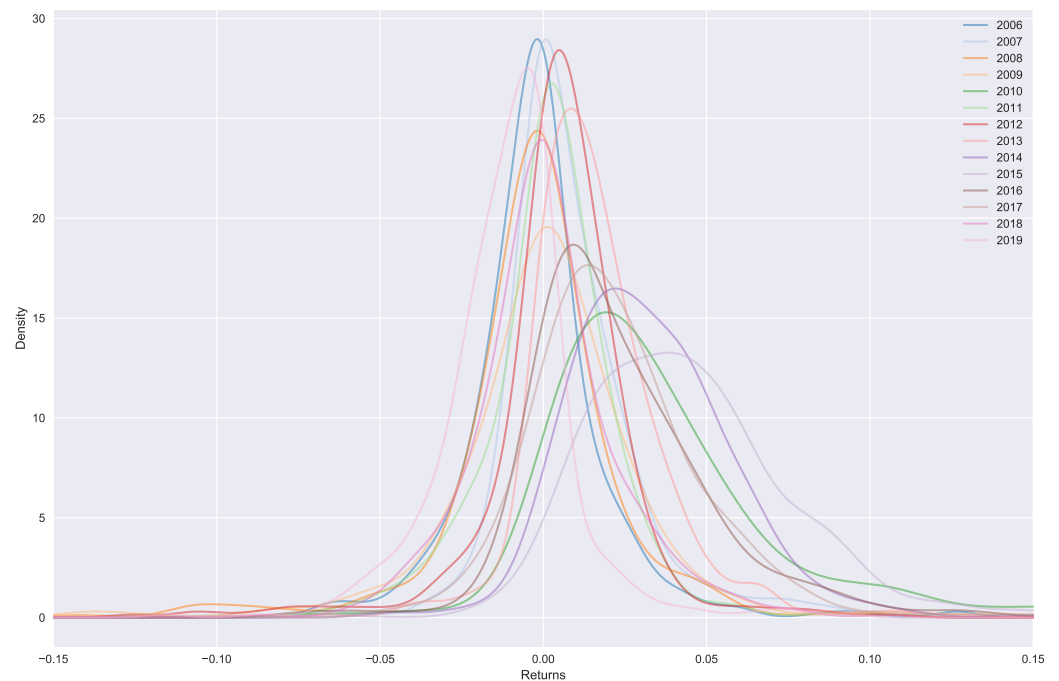


Figure A2. Density plots of returns grouped by year. We can see each year follows a different shape, but shows some striking regularities representing a statistical equilibrium.

Appendix C. Relation to Rational Inattention

In his seminal work, [2] outlined rational inattention “based on the idea that individual people have limited capacity for processing information”. This work introduced information-processing constraints into the macroeconomic literature, using mutual information as a measure of such information costs.

Of particular interest are the developments of [3] who showed how to apply rational inattention (RI) to discrete decision-making. The key contribution was the modification to the logit function that arises from considering a cost to decision-makers from deviating from prior knowledge. In this section, we highlight the similarities of R.I. with the thermodynamic approach of [4] and the work proposed here.

The problem to be solved is formulated as follows. A utility-maximising agent must make a discrete choice, while it is costly to acquire information about the options A available:

$$\max f[a, x] \sum_{a \in A} \int_x f[a, x] U[a, x] dx - T \left(- \sum_{a \in A} f[a, x] \log \left(\frac{f[a, x]}{p[x] f[a]} \right) \right) \tag{A6}$$

subject to $\sum_{a \in A} f[a|x] = 1$

where the first term is the expected utility, and the second a cost of information (following Sims [2], the mutual information). We see this as a similar setup to that of [4], which also corresponds to maximising the expected utility subject to an information cost, however, the information cost in [4] is instead measured as the KL-divergence. A key difference between the two is that Equation (A6) adds a dependence on $f[a]$ into the denominator of the information cost term. We can take the first order conditions of the resulting Lagrangian for (A6) and solve for $f[a|x]$, yielding:

$$f[a|x] = \frac{e^{\frac{U(a,x)}{T} + \log(f[a])}}{\sum_{a' \in A} e^{\frac{U(a',x)}{T} + \log(f[a'])}} = \frac{f[a] e^{\frac{U(a,x)}{T}}}{\sum_{a' \in A} f[a'] e^{\frac{U(a',x)}{T}}} \tag{A7}$$

which is not yet fully solved, as there is a dependence on the unconditional probability $f[a]$. Since $f[a] = \int_x f[a|x]p[x]dx$, $f[a]$ depends on $f[a|x]$, and $f[a|x]$ depends on $f[a]$, this must (generally) be solved numerically, for example, with the Blahut–Arimoto algorithm by first making a guess for $f[a]$ and then iterating from there (see Caplin et al. [73] or Matějka and McKay [3] for solutions). It is for this reason, we utilise the configuration of [4] for the decision-making component, which depends only on the prior probabilities, and not the unconditional action probabilities $f[a]$ meaning an analytical solution can be obtained. However, the R.I. framework can be seen as equivalent to choosing an “optimal” prior in the free energy framework of [4], as both can be seen as applications of rate-distortion theory [56].

Further discussion on the relationship between R.I. and QRSE is given in [30].

Appendix D. Additional Parameters

While μ and T are the main parameters of interest in this work, since they have a direct contribution to the modified decision function introduced, ρ and γ are still important, although to a lesser extent as they are indirectly impacted. ρ is the Lagrange multiplier for the competition constraint, and γ controls the skewness of the resulting distribution.

Appendix D.1. Impact of Decisions on Outcomes

Parameter ρ measures the impact of individual decisions on housing prices. A large ρ corresponds to a highly effective market (high impact of actions on the response). In contrast, a low ρ corresponds to a weaker market response, and thus, lower market effectiveness. Parameter ρ , therefore, corresponds to the strength of the negative feedback mechanism, with the case of $\rho = 0$ implying no market feedback (i.e., no impact on the outcome based on the actions). In all cases, we see relatively large ρ 's, peaking in 2013 and 2019, indicating the presence of a well-functioning feedback loop across the years. We see little variation between the uniform, previous, and mean prior in Figure A3, perhaps drawn from the fact the priors work as linear weightings in the difference between the conditional action probabilities, as shown in Equation (20).

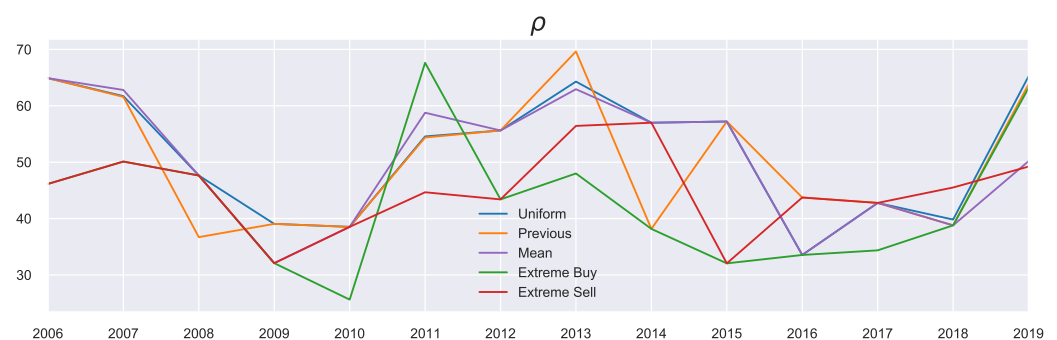


Figure A3. Competition.

Appendix D.2. Skewness

The parameter γ affects the skew of the resulting exponential distribution. This skew arises from (potentially) unfulfilled agent expectations, i.e., where $\mu \neq \xi$ [21]. Parameter γ , therefore, is a measure of skewness in the binary action case. In the asymmetric multi-action QRSE case, γ is replaced by alternate μ 's explaining such skew. As mentioned, the priors can also introduce such a skew (without the need for a γ). This is shown in the extreme buy γ in Figure A4 which was almost always near zero, as the buying preference already creates the skew needed to describe the underlying distribution (i.e., the skewness was already explained by p). In contrast, extreme sell needs small γ 's to switch their (incorrect) skew.

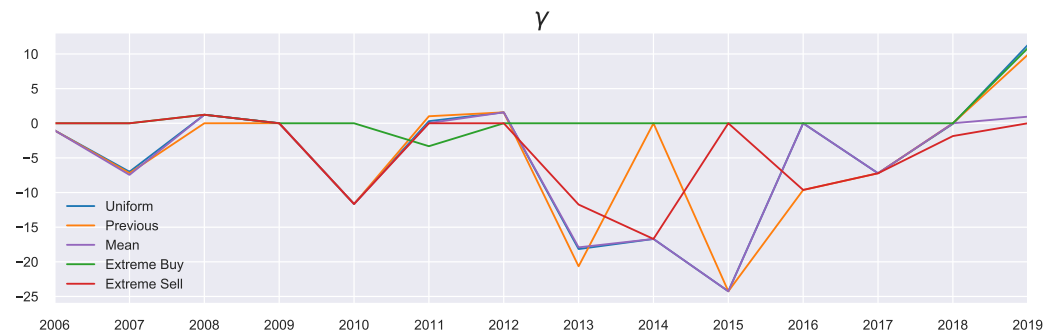


Figure A4. Skewness.

Negative γ corresponds to positive skewness, and positive γ corresponds to negative skewness. In most cases here, we see (at least slightly) positively skewed distributions (resulting in negative γ 's), with the exception of 2019, which is negatively skewed, as can be verified in Figure A5.

Generally, γ 's for the mean, previous, and uniform priors follow similar paths, except for the 2013–2016 years. In 2014 and 2016, γ 's for the previous priors differs from the other priors. This can be explained by the fact that in both cases, the prior had a strong sell preference (shown in Figure 4), meaning an adjusted γ was needed to capture the current distributions shift correctly (and offset the influence of the prior).

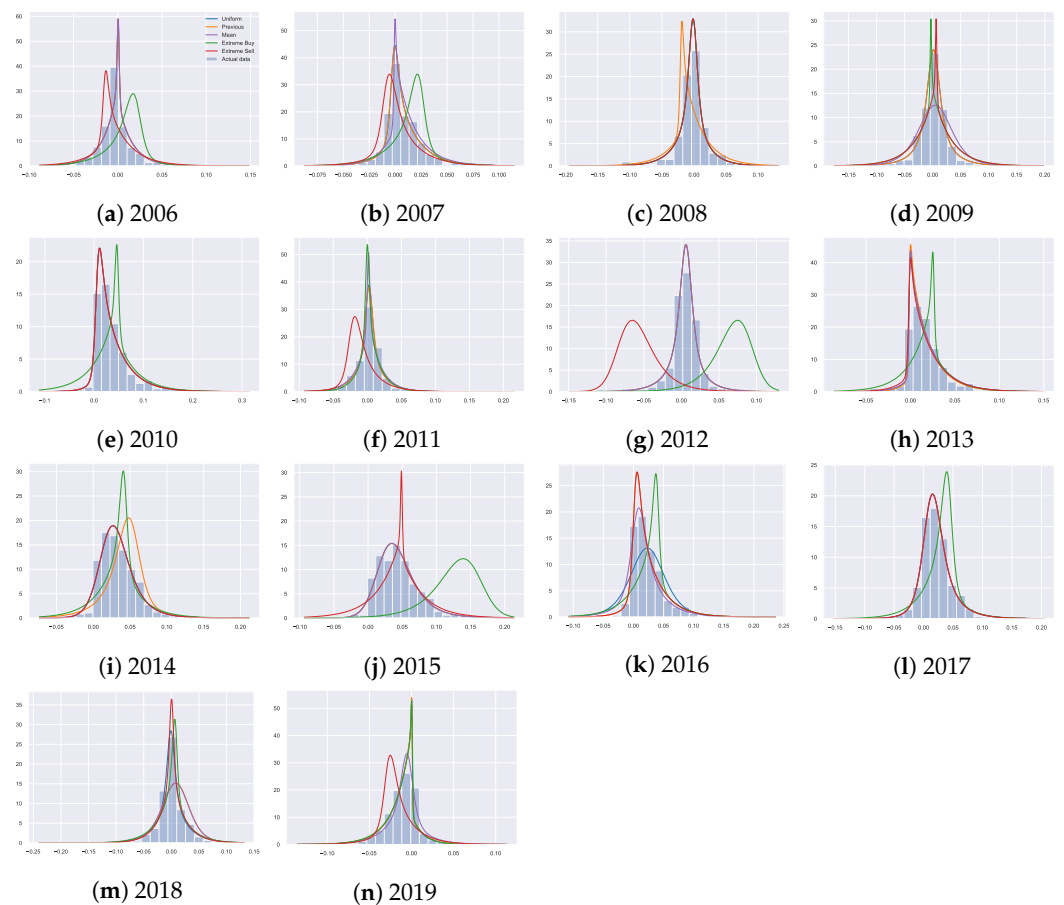


Figure A5. Resulting fitted marginals distributions $f[x]$ for each year. Each coloured line represents a different prior (with the legend given in the top left). The blue bars show the (discretized) actual return distribution.

Appendix E. Probability Plots

In this section, we provide the resulting probability plots for $f[x]$ (Figure A5), $f[a, x]$ (Figure A6), and $f[a|x]$ (Figure A7) across all years analysed.

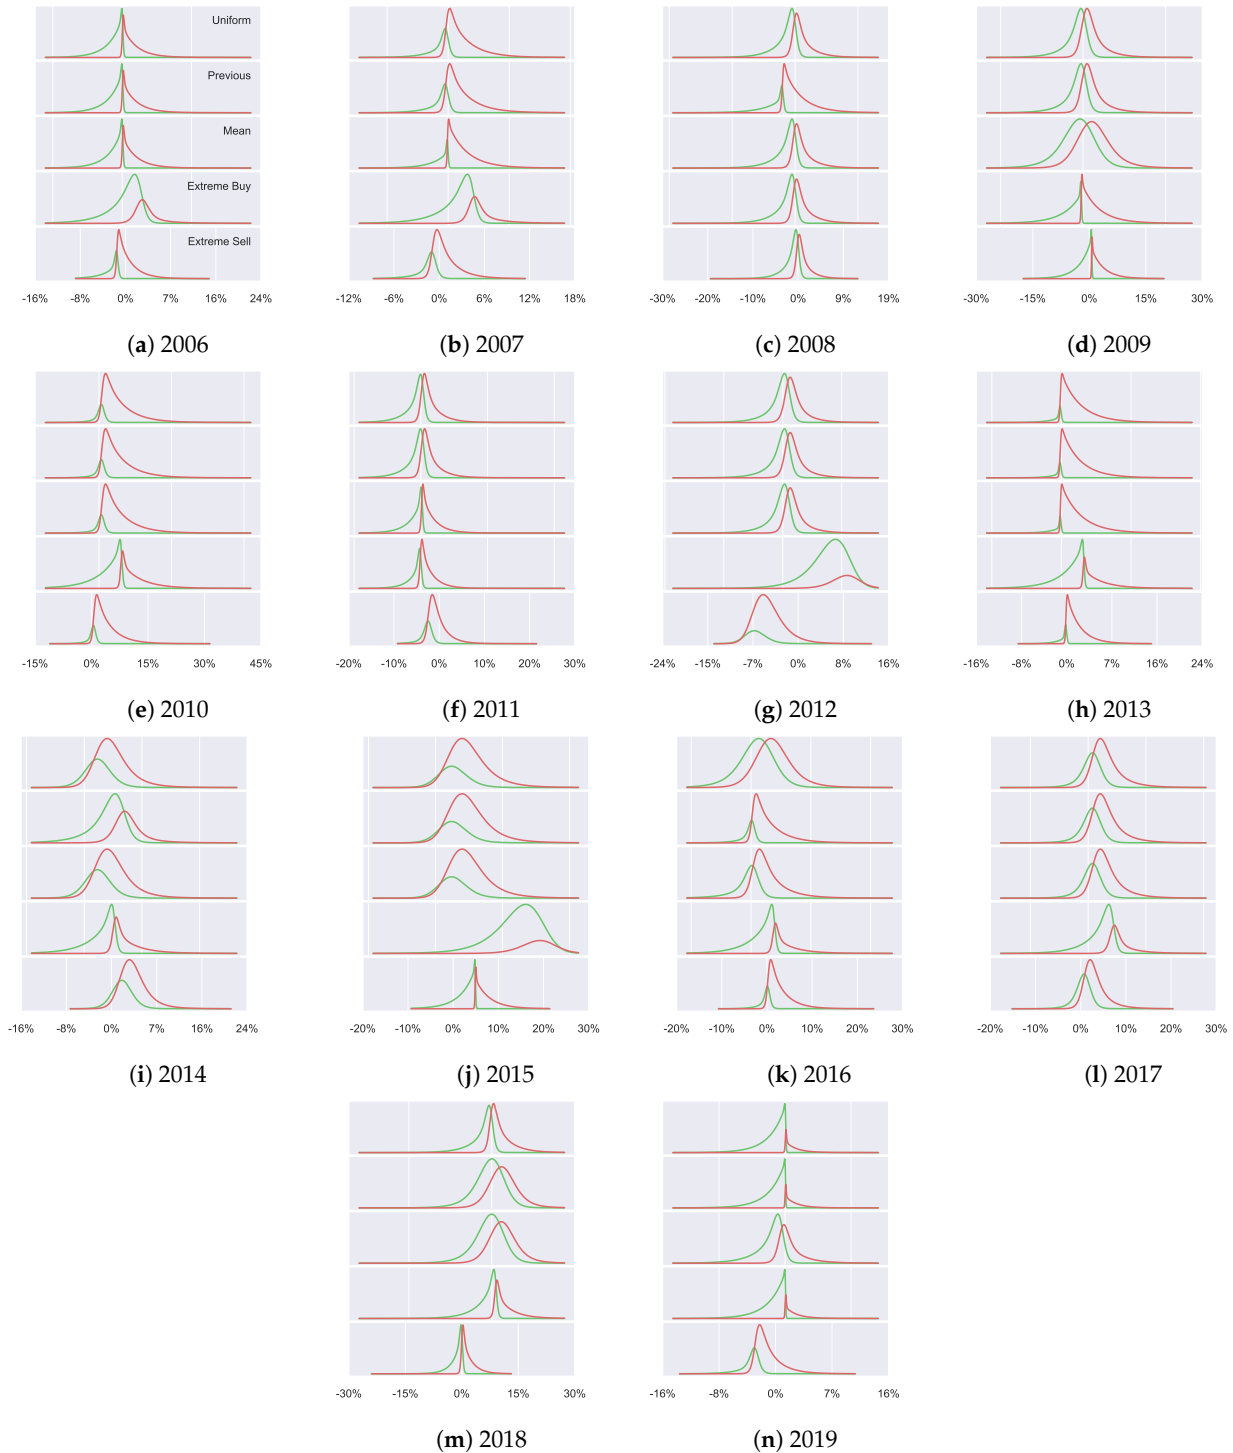


Figure A6. Resulting Joint Distributions. Red lines represent $f[\text{sell}, x]$, and green lines represent $f[\text{buy}, x]$. Each plot from top to bottom shows: Uniform, previous, mean and extreme buy and extreme sell priors (in that order).

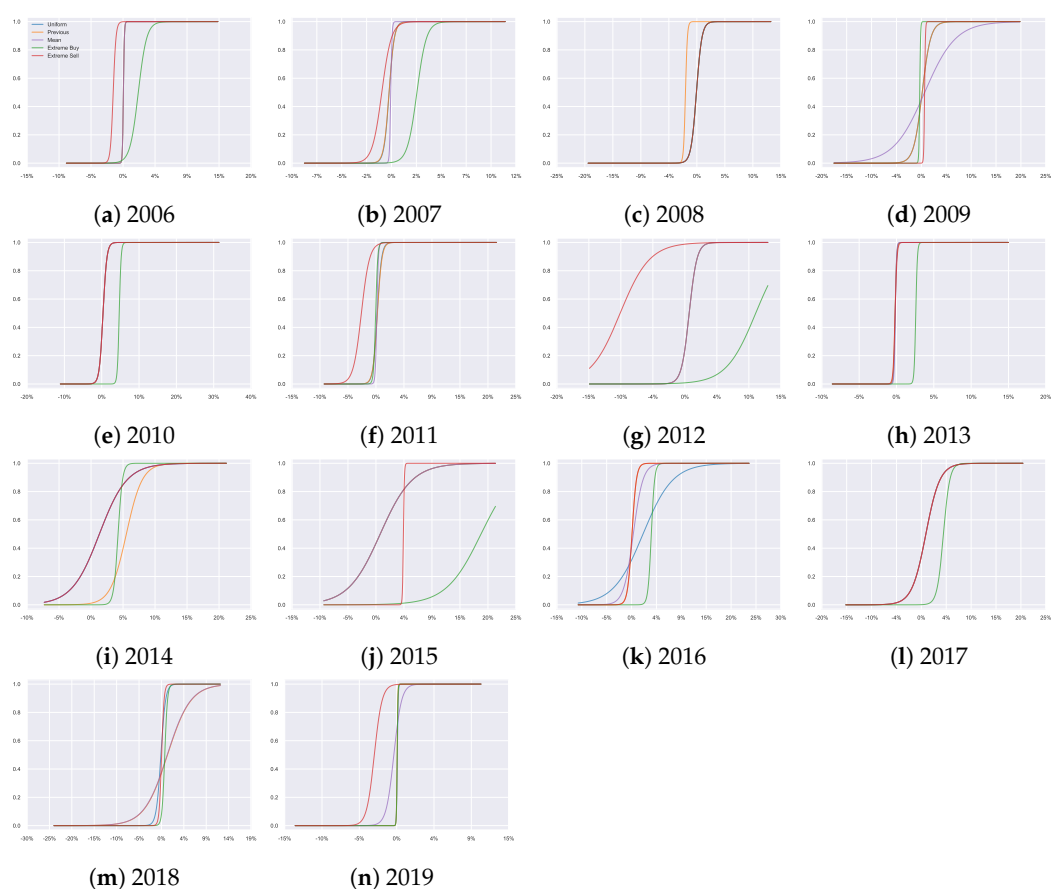


Figure A7. Decision functions for selling. Buying curves are excluded as they are simply the complement ($1 - \text{sell}$). The green lines represent the extreme buy a priori preference, which means the resulting probabilities of selling are shifted far to the right, i.e., the majority of the area comprises buying actions, and only the extreme positive growth rates for sell. In contrast, the red lines represent the sell preference, which “pulls” the area to the left, resulting in a strong resulting conditional preference for selling.

References

- Simon, H.A. *Models of Man; Social And Rational.*; Wiley: Hoboken, NJ, USA, 1957.
- Sims, C.A. Implications of rational inattention. *J. Monet. Econ.* **2003**, *50*, 665–690. [[CrossRef](#)]
- Matějka, F.; McKay, A. Rational inattention to discrete choices: A new foundation for the multinomial logit model. *Am. Econ. Rev.* **2015**, *105*, 272–298. [[CrossRef](#)]
- Ortega, P.A.; Braun, D.A. Thermodynamics as a theory of decision-making with information-processing costs. *Proc. R. Soc. A Math. Phys. Eng. Sci.* **2013**, *469*, 20120683. [[CrossRef](#)]
- Scharfenaker, E.; Foley, D.K. Quantal response statistical equilibrium in economic interactions: Theory and estimation. *Entropy* **2017**, *19*, 444. [[CrossRef](#)]
- Jaynes, E.T. Information theory and statistical mechanics. *Phys. Rev.* **1957**, *106*, 620. [[CrossRef](#)]
- Yang, J. Information theoretic approaches in economics. *J. Econ. Surv.* **2018**, *32*, 940–960. [[CrossRef](#)]
- Yakovenko, V.M. *Econophysics, Statistical Mechanics Approach to.* In *Encyclopedia of Complexity and Systems Science*; Meyers, R.A., Ed.; Springer: New York, NY, USA, 2009; pp. 2800–2826.
- Scharfenaker, E.; Semieniuk, G. A statistical equilibrium approach to the distribution of profit rates. *Metroeconomica* **2017**, *68*, 465–499. [[CrossRef](#)]
- Scharfenaker, E.; Yang, J. Maximum entropy economics. *Eur. Phys. J. Spec. Top.* **2020**, *229*, 1577–1590. [[CrossRef](#)]
- Wolpert, D.H.; Harré, M.; Olbrich, E.; Bertschinger, N.; Jost, J. Hysteresis effects of changing the parameters of noncooperative games. *Phys. Rev. E* **2012**, *85*, 036102. [[CrossRef](#)]
- Dragulescu, A.; Yakovenko, V.M. Statistical mechanics of money. *Eur. Phys. J. Condens. Matter Complex Syst.* **2000**, *17*, 723–729. [[CrossRef](#)]

13. Yakovenko, V.M.; Rosser Jr, J.B. Colloquium: Statistical mechanics of money, wealth, and income. *Rev. Mod. Phys.* **2009**, *81*, 1703. [CrossRef]
14. Foley, D.K. Unfulfilled Expectations: One Economist's History. In *Expectations*; Springer: Berlin/Heidelberg, Germany, 2020; pp. 3–17.
15. Harré, M.S. Information Theory for Agents in Artificial Intelligence, Psychology, and Economics. *Entropy* **2021**, *23*, 310. [CrossRef] [PubMed]
16. Foley, D.K. Information theory and behavior. *Eur. Phys. J. Spec. Top.* **2020**, *229*, 1591–1602. [CrossRef]
17. Ömer, Ö. Maximum entropy approach to market fluctuations as a promising alternative. *Eur. Phys. J. Spec. Top.* **2020**, *229*, 1715–1733. [CrossRef]
18. Yang, J.; Carro, A. Two tales of complex system analysis: MaxEnt and agent-based modeling. *Eur. Phys. J. Spec. Top.* **2020**, *229*, 1623–1643. [CrossRef]
19. Jaynes, E.T. Information theory and statistical mechanics. II. *Phys. Rev.* **1957**, *108*, 171. [CrossRef]
20. Jaynes, E.T. *Probability Theory: The Logic of Science*; Cambridge University Press: Cambridge, UK, 2003.
21. Scharfenaker, E. Implications of quantal response statistical equilibrium. *J. Econ. Dyn. Control.* **2020**, *119*, 103990. [CrossRef]
22. Ömer, Ö. Dynamics of the US Housing Market: A Quantal Response Statistical Equilibrium Approach. *Entropy* **2018**, *20*, 831. [CrossRef]
23. Ömer, Ö. Essays on Modeling Housing Markets, Income Distribution, and Wealth Concentration. Ph.D. Thesis, The New School, New York, NY, USA, 2018.
24. Ömer, Ö. Equilibrium-Disequilibrium Dynamics of the US Housing Market, 2000–2015: A Quantal Response Statistical Equilibrium Approach. Working Papers 1809, New School for Social Research, Department of Economics, 2018. Available online: <https://econpapers.repec.org/paper/newwpaper/1809.htm> (accessed on 30 September 2020)
25. Yang, J. A quantal response statistical equilibrium model of induced technical change in an interactive factor market: Firm-level evidence in the EU economies. *Entropy* **2018**, *20*, 156. [CrossRef]
26. Wiener, N. Measuring Labor Market Segmentation from Incomplete Data. Working Paper 2018-01, Amherst, MA, 2018. Available online: https://scholarworks.umass.edu/econ_workingpaper/238/ (accessed on 3 October 2020)
27. Wiener, N. Essays on Labor Mobility and Segmentation. Ph.D. Thesis, The New School, New York, NY, USA, 2019.
28. Wiener, N.M. Labor market segmentation and immigrant competition: A quantal response statistical equilibrium analysis. *Entropy* **2020**, *22*, 742. [CrossRef]
29. Blackwell, K. A Behavioral Foundation for Commonly Observed Distributions of Financial and Economic Data. Working Papers 1912, New School for Social Research, Department of Economics, 2019. Available online: <https://ideas.repec.org/p/new/wpaper/1912.html> (accessed on 8 October 2020)
30. Blackwell, K. Entropy Constrained Behavior in Financial Markets A Quantal Response Statistical Equilibrium Approach to Financial Modeling. Ph.D. Thesis, The New School, New York, NY, USA, 2018.
31. Scharfenaker, E. Statistical Equilibrium Methods in Analytical Political Economy. *J. Econ. Surv.* **2020**. [CrossRef]
32. Smith, A. *The Wealth of Nations: An inquiry into the nature and causes of the Wealth of Nations*; Harriman House Limited: Petersfield, UK, 2010.
33. McKelvey, R.D.; Palfrey, T.R. Quantal response equilibria for normal form games. *Games Econ. Behav.* **1995**, *10*, 6–38. [CrossRef]
34. McKelvey, R.D.; Palfrey, T.R. Quantal response equilibria for extensive form games. *Exp. Econ.* **1998**, *1*, 9–41. [CrossRef]
35. Lord, C.G.; Ross, L.; Lepper, M.R. Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *J. Personal. Soc. Psychol.* **1979**, *37*, 2098. [CrossRef]
36. K Levine, D. *Is Behavioral Economics Doomed?: The Ordinary Versus the Extraordinary*; Open Book Publishers: Cambridge, UK, 2012.
37. DellaVigna, S. Psychology and economics: Evidence from the field. *J. Econ. Lit.* **2009**, *47*, 315–372. [CrossRef]
38. Daunizeau, J.; Den Ouden, H.E.; Pessiglione, M.; Kiebel, S.J.; Stephan, K.E.; Friston, K.J. Observing the observer (I): Meta-bayesian models of learning and decision-making. *PLoS ONE* **2010**, *5*, e15554. [CrossRef] [PubMed]
39. Khalvati, K.; Park, S.A.; Mirbagheri, S.; Philippe, R.; Sestito, M.; Dreher, J.C.; Rao, R.P. Modeling other minds: Bayesian inference explains human choices in group decision-making. *Sci. Adv.* **2019**, *5*, eaax8783. [CrossRef] [PubMed]
40. Kruis, J.; Maris, G.; Marsman, M.; Bolsinova, M.; van der Maas, H.L. Deviations of rational choice: An integrative explanation of the endowment and several context effects. *Sci. Rep.* **2020**, *10*, 1–16. [CrossRef]
41. Debreu, G. Review of individual choice behavior by RD Luce. *Am. Econ. Rev.* **1960**, *50*, 186–188.
42. McFadden, D. Conditional logit analysis of qualitative choice behavior. In *Frontiers in Econometrics*, Zarembka, P., Ed.; Academic Press: Cambridge, MA, USA, 1973; pp. 105–142.
43. Golan, A. *Foundations of Info-Metrics: Modeling, Inference, and Imperfect Information*; Oxford University Press: Oxford, UK, 2018.
44. Hafner, D.; Ortega, P.A.; Ba, J.; Parr, T.; Friston, K.; Heess, N. Action and perception as divergence minimization. *arXiv* **2020**, arXiv:2009.01791.
45. Ortega, P.A.; Stocker, A.A. Human decision-making under limited time. In Proceedings of the Advances in Neural Information Processing Systems, Barcelona, Spain, 5–10 December 2016; pp. 100–108.
46. Gottwald, S.; Braun, D.A. The two kinds of free energy and the Bayesian revolution. *PLoS Comput. Biol.* **2020**, *16*, e1008420. [CrossRef] [PubMed]

47. Wilson, A. Boltzmann, Lotka and Volterra and spatial structural evolution: An integrated methodology for some dynamical systems. *J. R. Soc. Interface* **2008**, *5*, 865–871. [[CrossRef](#)]
48. Crosato, E.; Nigmatullin, R.; Prokopenko, M. On critical dynamics and thermodynamic efficiency of urban transformations. *R. Soc. Open Sci.* **2018**, *5*, 180863. [[CrossRef](#)]
49. Slavko, B.; Glavatskiy, K.; Prokopenko, M. Dynamic resettlement as a mechanism of phase transitions in urban configurations. *Phys. Rev. E* **2019**, *99*, 042143. [[CrossRef](#)]
50. Harding, N.; Spinney, R.E.; Prokopenko, M. Population mobility induced phase separation in SIS epidemic and social dynamics. *Sci. Rep.* **2020**, *10*, 1–11. [[CrossRef](#)]
51. Shore, J.; Johnson, R. Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy. *IEEE Trans. Inf. Theory* **1980**, *26*, 26–37. [[CrossRef](#)]
52. Kesavan, H.; Kapur, J. Maximum Entropy and Minimum Cross-Entropy Principles: Need for a Broader Perspective. In *Maximum Entropy and Bayesian Methods*; Springer: Berlin/Heidelberg, Germany, 1990; pp. 419–432.
53. Glavatskiy, K.S.; Prokopenko, M.; Carro, A.; Ormerod, P.; Harre, M. Explaining herding and volatility in the cyclical price dynamics of urban housing markets using a large-scale agent-based model. *SN Bus. Econ.* **2021**, *1*, 1–21. [[CrossRef](#)]
54. Evans, B.P.; Glavatskiy, K.; Harré, M.S.; Prokopenko, M. The impact of social influence in Australian real estate: Market forecasting with a spatial agent-based model. *J. Econ. Interact. Coord.* **2021**, 1–53.
55. Hommes, C.H. On the consistency of backward-looking expectations: The case of the cobweb. *J. Econ. Behav. Organ.* **1998**, *33*, 333–362. [[CrossRef](#)]
56. Genewein, T.; Leibfried, F.; Grau-Moya, J.; Braun, D.A. Bounded rationality, abstraction, and hierarchical decision-making: An information-theoretic optimality principle. *Front. Robot. AI* **2015**, *2*, 27. [[CrossRef](#)]
57. Friedman, M. *Theory of the Consumption Function*; Princeton University Press: Princeton, NJ, USA, 2018.
58. Hommes, C.; Wagener, F. Complex evolutionary systems in behavioral finance. In *Handbook of Financial Markets: Dynamics and Evolution*; Elsevier: Amsterdam, The Netherlands, 2009; pp. 217–276.
59. Evans, G.W.; Honkapohja, S. *Learning and Expectations in Macroeconomics*; Princeton University Press: Princeton, NJ, USA, 2012.
60. Chow, G.C. *Usefulness of Adaptive and Rational Expectations in Economics*; Center for Economic Policy Studies, Princeton University: Princeton, NJ, USA, 2011.
61. Soofi, E.S.; Retzer, J.J. Information indices: Unification and applications. *J. Econom.* **2002**, *107*, 17–40. [[CrossRef](#)]
62. Dziewulski, P. Just-noticeable difference as a behavioural foundation of the critical cost-efficiency index. *J. Econ. Theory* **2020**, *188*, 105071. [[CrossRef](#)]
63. Pfeifer, R.; Bongard, J. *How the Body Shapes the Way We Think: A New View of Intelligence*; MIT Press: Cambridge, MA, USA, 2006.
64. Polani, D.; Sporns, O.; Lungarella, M. How information and embodiment shape intelligent information processing. In *50 Years of Artificial Intelligence*; Springer: Berlin/Heidelberg, Germany, 2007; pp. 99–111.
65. Ay, N.; Bernigau, H.; Der, R.; Prokopenko, M. Information-driven self-organization: The dynamical system approach to autonomous robot behavior. *Theory Biosci.* **2012**, *131*, 161–179. [[CrossRef](#)]
66. Montúfar, G.; Ghazi-Zahedi, K.; Ay, N. A theory of cheap control in embodied systems. *PLoS Comput. Biol.* **2015**, *11*, e1004427. [[CrossRef](#)]
67. Polani, D.; Nehaniv, C.L.; Martinetz, T.; Kim, J.T. Relevant information in optimized persistence vs. progeny strategies. In Proceedings of the Artificial Life X: Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems, Bloomington, IN, USA, 3–6 June 2006; MIT Press: Cambridge, MA, USA, 2006.
68. Prokopenko, M.; Gerasimov, V.; Tanev, I. Measuring spatiotemporal coordination in a modular robotic system. In Proceedings of the Artificial Life X: Proceedings of the 10th International Conference on the Simulation and Synthesis of Living Systems, Bloomington, IN, USA, 3–6 June 2006; pp. 185–191.
69. Capdepuy, P.; Polani, D.; Nehaniv, C.L. Maximization of potential information flow as a universal utility for collective behaviour. In Proceedings of the 2007 IEEE Symposium on Artificial Life, Honolulu, HI, USA, 1–5 April 2007; pp. 207–213.
70. Tishby, N.; Polani, D. Information theory of decisions and actions. In *Perception-Action Cycle*; Springer: Berlin/Heidelberg, Germany, 2011; pp. 601–636.
71. Tishby, N.; Pereira, F.C.; Bialek, W. The information bottleneck method. *arXiv* **2000**, arXiv:physics/0004057.
72. Spinney, R.E.; Lizier, J.T.; Prokopenko, M. Transfer entropy in physical systems and the arrow of time. *Phys. Rev. E* **2016**, *94*, 022135. [[CrossRef](#)] [[PubMed](#)]
73. Caplin, A.; Dean, M.; Leahy, J. Rational inattention, optimal consideration sets, and stochastic choice. *Rev. Econ. Stud.* **2019**, *86*, 1061–1094. [[CrossRef](#)]