A Rational Inattention Unemployment Trap*

Martin Ellison[†] and Alistair Macaulay[‡]

May 6, 2019

Abstract

We show that introducing rational inattention into a model with uninsurable unemployment risk can generate multiple steady states, when the same model with full information has a unique steady state. The model features heterogeneity and persistence in household labour market expectations, consistent with survey evidence. In a heterogeneous agent New Keynesian model, we find that rational inattention to the future hiring rate generates a high employment steady state with moderate inflation, and an unemployment trap with very low (but positive) inflation and a low job hiring rate.

1 Introduction

There is a long history of papers suggesting that self-fulfilling expectations might allow an economy to become stuck in a bad steady state (see Diamond 1982 for an early example). One source of these fluctuations is the interaction of labour market expectations and precautionary saving. If households believe that their future employment prospects are bleak, they will increase precautionary savings today. The fall in aggregate demand that results causes employment to fall, confirming the pessimistic beliefs. This feedback loop is empirically important: Heathcote and Perri (2018) provide evidence that precautionary savings were a key driver of consumption around the onset of

*This paper supersedes an earlier paper titled 'Rational Inattention and Multiple Equilibria'. We thank Paul Beaudry, Gianluca Femminis, Albert Marcet, Filip Matějka, Edoardo Palombo, Morten Ravn, Ricardo Reis, Christopher Sims, Mirko Wiederholt, and conference and seminar participants at the Birkbeck Centre for Applied Macroeconomics - V: Annual Workshop, the China International Conference in Macroeconomics (CICM 2019) in Shenzhen, the Expectations in Dynamic Macroeconomic Models Workshop at the Barcelona GSE summer forum 2019, the MMF PhD Conference 2019 at City University of London, the New Advances on Informational Frictions in Business Cycles Workshop at the University of Lausanne, the 1st QMUL Economics and Finance Workshop for PhD & Post-doctoral students, the XII REDg Workshop in Quantitative Macroeconomics at the MOVE Foundation and Universitat Autonòma de Barcelona and the Institute of Economic Analysis-CSIC, the 2018 Royal Economic Society Annual Conference, Shanghai University of Finance and Economics, the University of Hong Kong, and the University of Oxford for helpful comments.

†Department of Economics & Nuffield College, University of Oxford and CEPR. Email: martin.ellison@economics.ox.ac.uk

‡Department of Economics & Nuffield College, University of Oxford. Email: alistair.macaulay@economics.ox.ac.uk

the Great Recession. Existing models of this mechanism have households precisely co-ordinating their labour market expectations on a particular equilibrium, as is common in models with multiple equilibria (Morris and Shin, 2000).

In this paper, we show that if it is costly for households to process information about future labour market conditions, their optimal information choices can generate multiple self-fulfilling steady states in a model which would have a unique steady state if households were fully informed. The unemployment trap generated in this way relies on the interaction of labour market expectations and precautionary saving, but it does not feature the strong co-ordination of household beliefs present in existing models of self-fulfilling labour market expectations.

There are two important assumptions that drive this result. Firstly, we assume that households cannot directly observe the hiring rate out of unemployment, which is crucial for precautionary saving in models with frictional labour markets (e.g. Ravn and Sterk, 2018). Households can obtain signals about the hiring rate, but this information processing is costly. Following Sims (2003) and others in the rational inattention literature, the cost is proportional to the informativeness of the signals, and so households choose to process somewhat noisy signals before deciding on their consumption. The hiring rate is naturally bounded between 0 and 1, so the commonly used result that rationally inattentive agents choose linear signals with Gaussian noise does not apply. The bounded support of the hiring rate implies that households choose signals with a discrete number of possible realizations even though the underlying variable is continuous (Matejka 2017). This nonlinear signal structure is what drives the possibility of multiple steady states.

The information processing cost implies that households have limited information about *realizations* of the hiring rate. Our second key assumption is that households also have limited information about the *structure* of their environment: they do not precisely know the true equilibrium marginal distribution of the hiring rate. This is related to the 'internal rationality' studied by Adam and Marcet (2011). To our knowledge, we are the first paper to examine this combination of rational inattention and imprecise prior beliefs. This is motivated by an observation from the Survey of Consumer Expectations: expectations have a much greater variance than the underlying hiring rate. This suggests that households do not have a good understanding of the range of values usually observed for the hiring rate.

The information choices made by households in this environment imply labour market expectations which are heterogeneous and persistent. There is also only weak co-ordination of beliefs: when the hiring rate falls, households on average receive signals that indicate it has fallen, and so average hiring rate expectations fall. These signals, however, are very noisy, so at the same time as the average expectation is falling a little, some households will be shifting their expectations up, and others will be expecting a large collapse in the hiring rate.

Survey labour market expectations also display these properties. In the Survey of Consumer Expectations, households are asked to predict the probability that, should they lose their main job today, they would find another suitable job within three months. This is precisely the hiring rate which drives precautionary saving in our model and in Ravn and Sterk (2018). Even after controlling for a wide range of personal characteristics, there is a great deal of disagreement about this rate each month, and expectations are highly persistent at the household level. We also find evidence that co-ordination of beliefs is weak, as in our model: if households all simultaneously agree on shifts in the hiring rate, then changes in average beliefs should account for a large amount of the variance of changes in household-level expectations over time. In fact, changes in average expectations account for just 0.1% of the variation in household belief changes. That is, most changes in household beliefs come from idiosyncratic factors. The noise in household signals in our model is a source of this kind of idiosyncratic variation.

The combination of information processing costs and imprecise prior beliefs is central to our results. When facing noisy signals, households bias their expectations towards their prior expectation. If households were to know before collecting any costly information that they are close to a particular steady state, their prior expectation would equal the true steady state hiring rate. Precautionary saving behaviour would therefore be very close to the optimal action under full information. Similarly, if households were to know when the economy has become stuck in an unemployment trap, and they were to know the hiring rate associated with that trap, then introducing rational inattention will have no effect on aggregate consumption. Since there are no non-linearities in our model except for those that arise endogenously as part of the optimal information choice, this would imply that the unemployment trap is not a steady state under full information or rational inattention.

Multiplicity does not disappear over time even if households observe a long history of signals and update their prior beliefs. In order to update their beliefs about the distribution of the hiring rate, households must process new information, which comes with a cost. Households process information until the marginal benefit equals this marginal cost. If prior beliefs already contain some information from previously observed signals, the optimal amount of new information processing will be very small. This logic is explored in detail by Matejka, Steiner, and Stewart (2017). Crucially, this prevents beliefs about the distribution of the hiring rate collapsing to the truth, which means that multiple steady states survive even when beliefs can update over time[1].

Section 2 places this work in the context of the literature. In section 3 we illustrate the potential for an unemployment trap in a simple static model. We then introduce our mechanism into a version of the HANK model from Ravn and Sterk (2018), in which the future hiring rate is very important in household decisions, in section 4. We show that the combination of rational inattention and imprecise prior beliefs about the hiring rate generates two possible steady states: a high employment steady state with a high hiring rate and moderate inflation, and an unemployment trap with a low hiring rate, and low (but positive) inflation. In section 5 we show that several key features of our model are found in survey data on hiring rate expectations. Section 6 concludes.

2 Related Literature

This paper relates to several strands of existing literature. Firstly, there is a vast literature studying fluctuations and traps driven by self-fulfilling expectations (see Cooper and John (1988) for a review of the early literature). Specifically, in our model changes in labour market expectations affect precautionary saving decisions, which in turn affect aggregate demand and so the labour market. Challe and Ragot (2016) show that a tractable model featuring this mechanism fits US data on aggregate consumption significantly better than benchmark models, and Challe et al (2017) find that the feedback loop between unemployment risk and precautionary saving played a significant role in the Great Recession. Beaudry et al (2017) show that it can lead to a 'Hayekian' recession after an over-accumulation of durable goods. Closely related to our paper are Heathcote and Perri (2018) and Ravn and Sterk (2018), who show that self-fulfilling labour market expectations can lead to the existence of multiple steady states: an economy can get stuck in a bad steady state where pessimistic beliefs persist indefinitely. We contribute to this literature by showing that these unemployment traps can be generated purely through the optimal information choices of rationally inattentive households, even in a model which is linear under full information. Moreover, our model does not require the co-ordination of beliefs which Morris and Shin (2000) argue is often necessary in models with multiple equilibria.

[1] In Adam and Marcet (2011), beliefs are 'near-rational', that is they cannot be empirically distinguished from the truth. The difference to our paper is that for Adam and Marcet, the only constraint on learning is the availability of data. For us, learning requires information processing, which is costly.

We also contribute to the literature on rational inattention. Most existing models with RI have agents with quadratic objective functions collecting costly information about a random variable with a known Gaussian distribution[2] (Sims, 2003, Mackowiak and Wiederholt, 2009). This has proved useful in explaining price stickiness (Mackowiak and Wiederholt, 2009, Matejka, 2015), consumption patterns (Sims, 2003, Luo, 2008, Luo et al, 2017), business cycle patterns (Mackowiak and Wiederholt, 2015), and other macroeconomic phenomena. Matejka (2017) and Jung et al (2015) show, however, that assuming a bounded prior belief leads to very different results to the quadratic-Gaussian formulation. Specifically, they show that the optimal decision rule of an agent facing a rational inattention problem with a bounded prior entails the agent choosing to limit themselves to a discrete number of options, even when the optimal choice under perfect information is continuous. As the probability of finding a job is naturally bounded by 0 and 1, our model displays these features. This paper is therefore a response to the call in Sims (2006) to explore the implications of RI away from the quadratic-Gaussian case, and to our knowledge we are the first to incorporate RI with bounded prior beliefs into a general equilibrium setting.

Our framework also relates closely to the literature on internal rationality. Adam and Marcet (2011) show that allowing for internally rational agents, who optimise given their beliefs but do not precisely know the equilibrium distributions of state variables, has important effects on asset pricing models. Adam et al (2012) use this to explain movements in house prices and the current account. We extend this literature by showing that the interaction of internal rationality and information processing costs creates multiplicity, where neither assumption generates this by itself. We believe that we are the first to combine these two information restrictions. We deviate from existing literature on internal rationality in that we do not assume that agent beliefs are close to the true equilibrium distribution of the state variable. The typical logic for 'near-rationality' is that agents learn over time, so would discard any beliefs which are very dissimilar to the truth. This does not happen when households learn in our model, because to learn they must process costly information, which means households stop learning before they reach near-rational beliefs.

Finally, our work contributes to the literature on heterogeneous expectations in macroeconomics. Armantier et al (2015) and Meeks and Monti (2018) show that inflation expectations are heterogeneous across households; in section 5 we document that the same is true for hiring rate expectations. The theoretical implications of heterogeneity in inflation expectations have been

[2]This is convenient as the optimal posterior belief about the shock, after processing information, is also Gaussian.

studied by Andrade et al (2019), Wiederholt (2017), among others. In contrast to this literature, we study heterogeneous labour market expectations.

3 Static Model

This section illustrates our mechanism in a simple static model. There are two sub-periods in this setup, which we refer to as morning and afternoon. The consumer problem is related to Ravn and Sterk (2018) and to our HANK model in section 4, but is substantially simplified by the static nature of the problem. The firm side is also kept very simple, to illustrate the key forces driving the unemployment trap. In section 3.8 we construct a dynamic model by repeating this static model, allowing households to update their prior beliefs over time using information processed in previous periods. This demonstrates that the unemployment trap does not disappear even when households update their prior beliefs over time.

3.1 Households

There is a unit mass of households. All households are employed in the morning, and they receive a wage of 1. With probability ω a household loses their job at the end of the morning. There is then a round of hiring by firms, so newly separated workers find employment for the afternoon with probability η. If they are employed in the afternoon, the household receives a wage w. If they are unemployed they receive the value of home production $\theta < w$.

Households can save or borrow in the morning at interest rate R. The budget constraints for household i in the morning, afternoon if employed, and afternoon if unemployed (respectively) are:

$$c_{mi} + \frac{b_{mi}}{R} = 1 \tag{1}$$

$$c_{ai}^{e} = w + b_{mi} \tag{2}$$

$$c_{ai}^{u} = \theta + b_{mi} \tag{3}$$

Combining these we have:

$$c_{ai}^{e} = w + R(1 - c_{mi}) \tag{4}$$

$$c_{ai}^{u} = \theta + R(1 - c_{mi}) \tag{5}$$

Households have quadratic consumption utility: $U(c_{ti}) = -\frac{1}{2}(c - \bar{c})^2$.

The household problem is therefore to choose morning consumption c_{mi} to maximise:

$$V_{mi} = -\frac{1}{2}(c_{mi} - \bar{c})^2 - \frac{1}{2}\beta \mathbf{E}_{mi}\omega(1-\eta)(c^u_{ai} - \bar{c})^2 - \frac{1}{2}\beta \mathbf{E}_{mi}(1 - \omega(1-\eta))(c^e_{ai} - \bar{c})^2 \tag{6}$$

If the households observe the hiring rate η all households make the same choice of morning consumption c_{mi}, to satisfy the FOC (dropping the i subscripts):

$$c_m = \beta R\omega(1-\eta)\big(\theta + R(1-c_m)\big) + \beta R\big(1 - \omega(1-\eta)\big)\big(w + R(1-c_m)\big) \tag{7}$$

In this case, morning consumption is an increasing linear function of the afternoon hiring rate, due to a simple precautionary savings motive. Solving out for morning consumption we have:

$$c_m = \frac{\beta R}{1+\beta R^2}\big(R + w - \omega(1-\eta)(w-\theta)\big) \tag{8}$$

3.2 Rational Inattention Problem

We now relax the assumption that households can precisely observe the afternoon hiring rate η. Instead, we assume that they can collect information about η from a variety of noisy signals, but doing so is costly. This cost is increasing in the informativeness of the signal chosen. This is formalised in equation 12 below. As well as the amount of information in the signal (which we will denote κ), the agent must also choose how this information is to be structured: they could choose a signal which is very accurate in some ranges of η but not in others, for example.

The payoff function being maximised is as in equation 6. The agent views the hiring rate as exogenous[3]. In maximising their expected utility the agent must decide on an optimal decision rule to map the signals they are able to process to consumption.

The solution to this problem therefore takes the form of an *information strategy* and an *action strategy.* The information strategy gives the amount of information the agent should process, and what form the signals should take. The action strategy maps from signal realizations to consumption choices.

[3]The hiring rate will in fact be endogenous to aggregate agent choices, but the agent does not take this into account. This will be explored in more detail in section 3.7. Notice that this assumption means that higher order beliefs do not affect the household problem.

This household problem is closely related to the firm profit maximisation problem in Matejka (2017). We therefore proceed by following the steps used to solve Matejka's firm problem. First, we simplify the household problem by noting that the value function in equation 6 is strictly concave in morning consumption c_{mi}, so there is a unique function mapping the optimal morning consumption choice c^*_{mi} to the expectation of the afternoon hiring rate formed by the household in the morning, $\mathbf{E}_{mi}\eta$. Furthermore, the optimal morning consumption choice is a continuous and strictly increasing function of the expected hiring rate, so there is a one-to-one mapping between the expected afternoon hiring rate and optimal morning consumption. A household will never choose a signal structure that has two distinct possible realizations that imply the same expected hiring rate, because distinguishing between the two realizations is a waste of information processing. There will therefore be a one-to-one mapping between signal realizations and the optimal morning consumption choice. We can therefore leave the signal choice in the background of the problem, and instead study the optimal decision rule linking morning consumption to the hiring rate, subject to the information costs of implementing such a rule.

Specifically, we express the household's decision rule as $f_i(\eta, c_{mi})$, the joint probability density function over the hiring rate and morning consumption. That is, given a particular afternoon hiring rate η, the household chooses how often they will choose each different possible level of morning consumption. They are aware that signals contain noise, so they are deciding how often, and by how much, they are willing to choose the wrong c_{mi} for each level of η, given the information costs of reducing those mistakes.

The problem of household i is therefore:

$$f_i = \arg\max_{\hat{f}_i} \mathbf{E}_{mi}[V(\eta, c_{mi})] - \psi\kappa(\hat{f}_i, g_i) = \arg\max_{\hat{f}_i} \int_\eta \int_{c_{mi}} V(\eta, c_{mi})\hat{f}_i(\eta, c_{mi})d\eta dc_{mi} - \psi\kappa(\hat{f}_i, g_i) \tag{9}$$

subject to

$$\int_{c_{mi}} \hat{f}_i(\eta, c_{mi})dc_{mi} = g_i(\eta) \qquad \forall\eta \tag{10}$$

$$\hat{f}_i(\eta, c_{mi}) \geq 0 \qquad \forall\eta, c_{mi} \tag{11}$$

$$\kappa(\hat{f}_i, g_i) = H[g(\eta)] - \mathbf{E}_{c_{mi}} H[\hat{f}_i(\eta|c_{mi})] \tag{12}$$

The function $H[.]$ is the entropy of the distribution over which it operates. That is:

$$H[g_i(\eta)] = -\int g_i(\eta)\log g_i(\eta)d\eta \tag{13}$$

The first constraint (equation 10) ensures that the marginal distribution of the hiring rate η obtained from the optimal joint pdf is consistent with $g_i(\eta)$, household i's prior belief about the distribution of the hiring rate [4].

The second constraint (equation 11) is that the solution must be positive everywhere, as required for the decision rule to be a joint pdf.

The final constraint (equation 12) is the information processing constraint. Entropy $H[.]$ is a measure of the dispersion of a distribution. The first term of constraint 12 is the entropy of the prior. The prior reflects the information held by the household about the distribution of the hiring rate before receiving any signals. We will assume for now that this prior is the same for all households and is uniformly distributed (this is relaxed in section 3.8), so the prior is rather dispersed and entropy is high. The second term is the expected entropy of $f_i(\eta|c_{mi})$, the updated distribution over the hiring rate believed by the household after taking in the available signals[5]. A precise posterior knowledge of η would give a very low posterior entropy, so the entropy difference from the prior would be large. Information costs in this model are proportional to this difference, that is how much the agent can learn from the signals. Note that with identical prior beliefs, each household now faces exactly the same problem, and so each household chooses the same decision rule f_i. We therefore drop the i subscripts from the decision rule and the prior belief g. Households will still choose heterogeneous values for morning consumption, because the optimal signals contain idiosyncratic noise.

Given a particular level of information processing κ, the household must decide how to allocate that 'information processing budget'. They could, for example, ensure they make no mistakes at all when the hiring rate is above a certain threshold, but in doing so they must accept that they will make larger mistakes with higher probability when η is below that level, or pay for more

[4]In Matejka (2017), this marginal distribution of the variable subject to rational inattention is the true distribution of that variable. This will not be the case here, as η will be determined endogenously in the model. Instead, the marginal distribution obtained by integrating the joint pdf over consumption should be interpreted here as the distribution of the hiring rate the household is expecting to see from their 'ignorance prior' (see section 3.7).

[5]The 'information content' of the signals is incorporated into the choice of c_{mi}, so the conditional distribution of η given the choice of c_{mi} tells us what the agent believes about η. This is because of the one-to-one mapping between signals and actions discussed above.

processing capacity. This means that the household faces a trade-off: for a particular κ they can distinguish between a several values of η which are close together, but that reduces the entropy of the posterior a great deal, so outside of that small range of η their posterior $f(\eta|c_{mi})$ must remain dispersed. When η is in that small range, the household making that decision will be very accurate in choosing optimal c_{mi}, but when η is outside of that range they will make large mistakes with a high probability. Alternatively, they can choose to allocate their information processing capacity to distinguishing between a small number of cases which are far apart. They are then never precise in predicting the hiring rate, but they make large mistakes less often. This is what drives the result in Matejka (2017) that the agent optimally restricts themselves to a small number of discrete levels of the choice variable, even though a continuous range of that variable is available, when the optimal κ is sufficiently small that the information constraint binds.

3.3 Rational Inattention Solution

The hiring rate is naturally bounded by 0 and 1. Assume that prior beliefs are uniform over the whole of this range: $g(\eta) \sim [0, 1]$. The optimal decision rule for a marginal cost of information of $\psi = 0.002$ is plotted in figure 1. This cost is sufficiently high that the optimal information strategy is to collect signals which are less than perfectly informative about η. The information processed at this cost is such that the household optimally chooses to restrict themselves to two levels of consumption, even though under perfect information the optimal consumption choice is continuous in the hiring rate. This is similar to the firm problem studied in Matejka (2017). The logic behind this is discussed in section 3.2 above, and in detail in Jung et al (2015) and Matejka (2017). As the hiring rate increases (and so the optimal choice of morning consumption under perfect information increases), the probability a household chooses the higher level of morning consumption in their 'menu' increases.

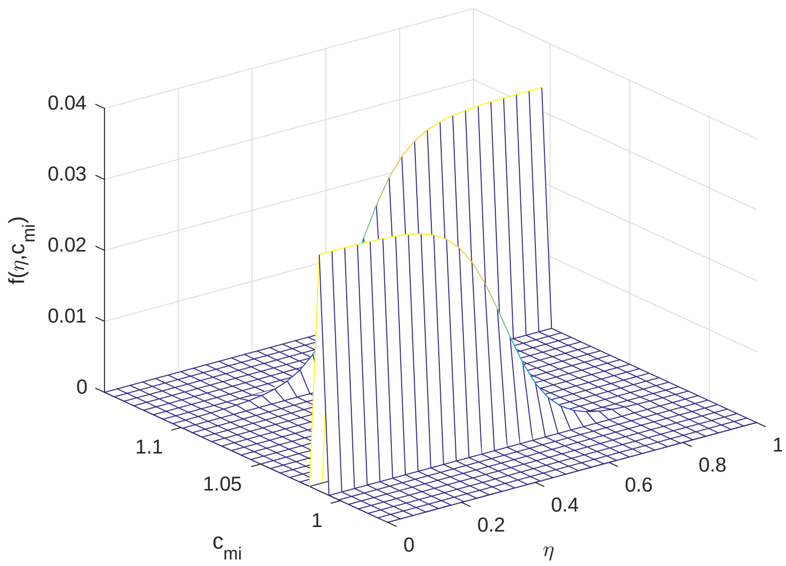

Figure 1: Optimal decision rule for $\psi = 0.002$

3.4 Aggregate Consumption

There is a unit mass of households making this decision. They all face the same labour market conditions, but we assume that they receive different noisy signals and/or they interpret those signals differently. Therefore for each level of the hiring rate some agents choose each of the morning consumption levels in the optimal 'menu' which arises from the RI problem with uniform priors (equations 9 - 12). The proportions on each level of consumption are determined by the probabilities in the optimal joint pdf obtained as the decision rule from the household problem.

Therefore for each level of η we obtain aggregate morning consumption c_m using:

$$c_m(\eta) = \int_{-\infty}^{\infty} c_{mi} f(c_{mi}|\eta) dc_{mi} \tag{14}$$

With no information processing[6], households choose consumption to maximise $\mathbf{E}_{mi}V$ given their prior belief (where $\mathbf{E}_{mi}\eta = 0.5$ for all households), as they cannot update their beliefs beyond that. They do not therefore change their consumption choice at all as the hiring rate varies. As the optimal information processing capacity κ increases (as ψ decreases), some information about η is processed, so households begin to choose different levels of c_{mi} for different underlying η. For low values of κ, households optimally restrict themselves to two values of morning consumption. Importantly, the aggregate morning consumption function has a wave-like shape around its perfect

[6] Households cannot process negative amounts of information - they cannot choose to forget information in their prior belief. That means that for any information cost ψ above a certain threshold, households hit this no-forgetting constraint and choose to process precisely no information ($\kappa = 0$).

information equivalent.

As shown in Matejka (2017), as information processing κ rises further, more choices of c_{mi} are introduced into the optimal menu. As this occurs the aggregate response of morning consumption to the hiring rate approaches the perfect information first order condition. For ease of exposition, the graphs in this paper are all drawn for information costs that imply households choose an optimal menu with two levels of consumption, but this is not important for the results. An example with a lower information cost in this static model is shown in appendix A.

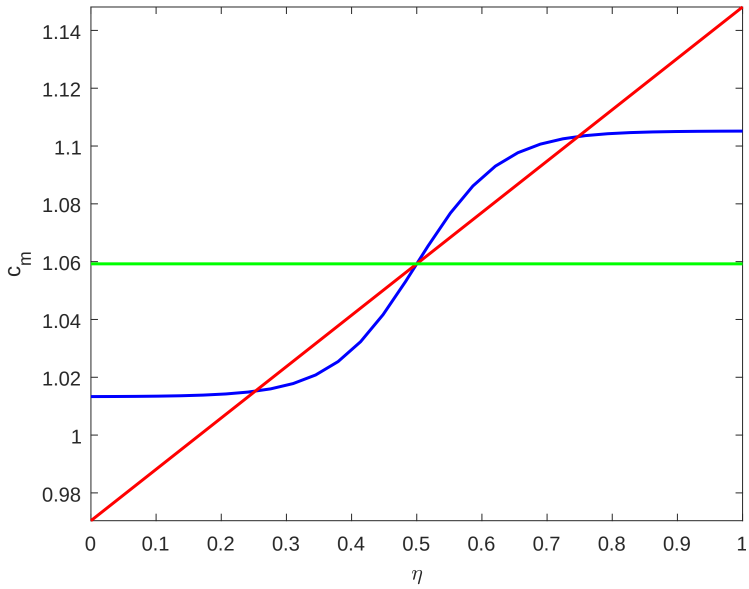

Figure 2: Aggregate morning consumption function with $\kappa = 0$ (green), $\kappa = 0.5$ (blue) and in the unconstrained case (red)

Consider the case where information processing is constrained but non-zero (the blue curve). This meets the full information consumption function (in red) at $c_m = 1.06, \eta = 0.5$. However, at this hiring rate in the full information model, every household consumes the same amount. In contrast, in the rational inattention model, half of the households get a signal that the hiring rate is 'high' and choose the high level of $c_{mi} = 1.104$. The other half receive a signal that η is 'low', and so consume the lower level $c_{mi} = 1.014$.

The flat sections of the aggregate consumption function under rational inattention occur where changes in the hiring rate do not lead to much change in the proportions of agents choosing each level of consumption in their menus. In figure 1 above, it can be seen that this is the case for extreme high and low values of η, and $c_m(\eta)$ is flat in these regions accordingly. In contrast, as η moves from 0.4 to 0.6, large numbers of households switch from choosing the low level of consumption to the high level, and this corresponds to the steep section of the corresponding aggregate morning consumption function in figure 2.

The shape of the aggregate consumption function is therefore driven by the shape of the curves in the optimal decision rule: if the probability of choosing the low level of consumption in figure 1 fell linearly as η increased the aggregate consumption function would be linear. In fact, the distribution of $c_{mi}|\eta$ for the values of c_{mi} in the optimal menu is logistic in shape, which is what gives rise to the wave-like shape of the aggregate response curve[7]. Beliefs, and so choices, are therefore endogenously sticky in certain regions of the support of η, as a result of the optimal signal structure that comes out of the entropy-based cost function. It is this non-linearity which leads to multiple equilibria in our model.

3.5 Firms

Usually, rational inattention models specify that agents collect costly information about exogenous variables, often shock processes. In contrast, we assume that the hiring rate is determined in equilibrium by firm hiring decisions[8]. For the purposes of this simple static example, it will be sufficient to say that firms hire more workers when aggregate demand is high, so:

$$\eta = H(c_m), \text{ with } H'(c_m) > 0 \tag{15}$$

The focus of this static model is household information choices, and the aggregate consumption function these imply, so for simplicity we use a linear H function throughout this section, though this is not necessary for our results. Appendix B microfounds such a process for the hiring rate using a model with working capital. In section 4 the firm side of the model is standard, as in Ravn and Sterk (2018).

3.6 Equilibrium

It is important here that the hiring rate η is not necessarily uniformly distributed like the prior beliefs of the households. In Matejka (2017) and many other models in the rational inattention literature, agents' prior beliefs about the distribution of the unknown variable are correct. We extend the RI literature by considering prior beliefs which do not precisely match the true equi-

[7]Matejka and McKay (2015) study in detail how RI leads to the logit model.

[8]In a standard model where agents perfectly understand how endogenous variables are determined this distinction is irrelevant, as agents know the mapping from shocks to endogenous variables. The distinction is relevant here because we assume that households do not know how the hiring rate is determined in equilibrium.

librium distribution of the relevant variables. This may be a plausible situation for information processing about variables which are difficult to learn about, perhaps because they are not easily understood, or because the data is not reported at the front of central bank communications and other news sources, so the variable's history is not easy to observe. This assumption is related to the 'internal rationality' in Adam and Marcet (2011). The survey data in section 5 suggests that for the hiring rate, households do not know the true equilibrium distribution.

In particular, we begin by assuming that households have a uniform 'ignorance' prior, which would be justified if the households do not understand how their decisions (and those of other agents) affect the hiring rate[9]. This is discussed in section 3.7. We study equilibrium in the static model with these uniform priors. In section 3.8 we take this as a starting point and repeat the static problem many times, allowing households to update their prior beliefs each period. We show that the multiplicity does not disappear when agents update their priors away from the uniform starting point.

The separation of the prior belief and the true equilibrium distribution of η means that the aggregate consumption function $c_m(\eta)$ remains as in figure 2, and the hiring rate is then determined endogenously by the interaction of this aggregate consumption and the firm hiring function 15.

The graph below shows this equilibrium interaction. The blue and red aggregate consumption functions are as in figure 2. The equilibrium condition 15 is added in black. Under full information there is one equilibrium, but under RI there are two: a high employment equilibrium and an unemployment trap with low morning consumption, and a low hiring rate, so low employment.

[9]The initial prior does not need to be necessarily uniform for our results. We require a prior belief which is significantly dispersed, hence the term 'ignorance prior', but the precise distribution is unimportant. The prior belief must also be bounded, which is ensured in this case as the hiring rate must be between 0 and 1 by definition.

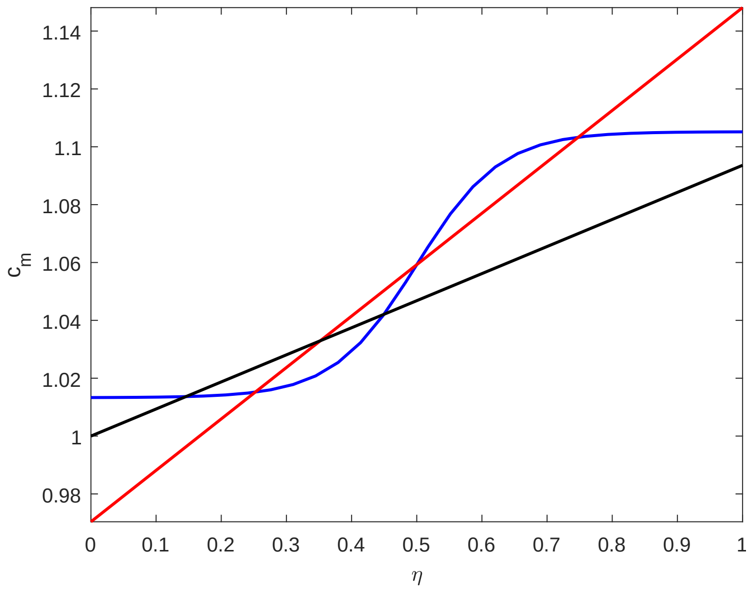

Figure 3: Aggregate consumption response to changes in the hiring rate for full information (red) and rational inattention with $\psi = 0.002$ (blue), with the firm condition 15 in black

Consider first the high employment equilibrium under rational inattention, at $\eta = 0.44$. As explained in section 3.4, at this hiring rate, close to the middle of the uniform prior beliefs, aggregate consumption is close to that under full information, but there is dispersion in household choices underlying this which is not present with full information. The key difference between aggregate consumption in these two models comes when η falls a little from this central equilibrium. Under perfect information, all agents respond the same way, by reducing their consumption a little. With rational inattention, however, the fall in the hiring rate leads to a large number of agents switching from the high level of morning consumption $c_{mi} = 1.104$ to the low level $c_{mi} = 1.014$. This means that aggregate consumption falls a great deal, so much so that there is another equilibrium at $(\eta = 0.17, c_m = 1.016)$ in which almost all households choose the low level of consumption. This is the unemployment trap.

In effect, households are using their limited information processing capacity to decide if they face a 'high' or 'low' hiring rate. As η moves a little below 0.5, the majority of agents decide on 'low', and consume accordingly, whereas if they knew the hiring rate more precisely they would choose consumption based on only a 'somewhat low' η. However that in itself would not be sufficient for multiplicity. It is crucial that, as the hiring rate falls even further, households do not decide on even lower consumption. They continue to believe that the hiring rate is 'low' even when it become 'extremely low'. That is why the aggregate consumption function flattens out at very low values of the hiring rate, which is why there is an equilibrium with very little labour market activity.

3.7 Requirements for multiplicity

The multiplicity of equilibria is driven by the non-linearity in the aggregate consumption function, which arises endogenously from the optimal information processing decisions of households. The first key requirement for our result is that households process costly information about the hiring rate, with an 'ignorance prior'. That is, households have only imprecise knowledge of the true equilibrium distribution of the hiring rate.

To see why, consider the equilibrium in figure 3. The true equilibrium distribution of η contains just two discrete points. If household prior beliefs matched this distribution, the optimal level of consumption in their decision rule when they received a signal that η was low would be closer to 1.06 (the optimal value when $\eta = 0.5$), as they do not need to worry about the possibility of an extremely low η. This, in turn, would mean that the equilibrium value of η in the unemployment trap would be higher. Iterating this logic, we can see that if beliefs were indeed close to the true distribution of η, the only equilibrium to survive would be the full information equilibrium.

The households in most standard rational inattention papers (e.g. Mackowiak and Wiederholt, 2015) do not have ignorance priors. In such models households understand all of the direct and general equilibrium effects mapping shocks into endogenous variables, and they know the distribution of shocks, so they can deduce the equilibrium distributions of all endogenous variables. That kind of model therefore assumes that households have a large amount of information about mechanisms at work in the economy (including how other households make decisions), but then it places limits on the information that those very well-informed households can obtain about the *realisations* of the shocks.

Our assumption of 'ignorance priors' builds on these existing rational inattention models by assuming that households do not perfectly understand the links from shocks to aggregate household choices to the hiring rate[10]. This assumption is related to the removal of 'external rationality' in Adam and Marcet (2011). Like them, we believe that it is plausible that agents do not precisely know the true stochastic processes and mechanisms underlying the determination of the endogenous variables they face, especially when those mechanisms are complicated and involve the choice behaviour of many other agents[11]. The survey data on hiring rate expectations in section 5 sup-

[10]This is significantly easier to solve than a general equilibrium rational inattention model without ignorance priors, because higher order beliefs cannot matter for household choices when households do not understand how the actions of others feed into the hiring rate

[11]We also follow Adam and Marcet in assuming that the agents in our model make choices that rationally maximise expected utility given their beliefs.

ports this view: the variance in survey expectations (even after accounting for a wide variety of household characteristics) is an order of magnitude larger than the long run variance of the true hiring rate. It appears that, just as in our model, households do not know the appropriate range of values for the hiring rate.

However, even if households do not understand the mechanisms leading to a particular equilibrium distribution of the hiring rate, they could obtain accurate prior beliefs if they had observed the hiring rate over many periods, and had learned its equilibrium distribution over time. This is related to the concept of 'near-rationality' in Adam and Marcet (2011): in their model subjective beliefs are sufficiently close to the true equilibrium process that agents cannot distinguish between the two given the data they observe. In contrast, we start from a uniform prior, which is extremely dispersed and far from the true equilibrium distribution. In fact, we do not require that the prior is uniform, but it does need to have a high degree of dispersion and be bounded. In section 3.8 we show that as long as agents start at some point in time with such dispersed priors, and they can change their information processing over time, our results hold even when households update their prior beliefs each period. This is because the initial reduction in dispersion of prior beliefs after one period of information processing leads households to process less information in the following periods, as households then rely more on their more informative priors and less on new signals (as in Matejka, Steiner, and Stewart, 2017). Prior beliefs therefore do not approach the true distribution of the endogenous variable[12].

The other important condition for multiplicity is that the firm hiring function $H(c_m)$ must be upward sloping. The aggregate consumption function is non-linear, but is always upward sloping. If it was downward sloping at any point, then an increase in the hiring rate must be leading to a rise in the probability that households get a signal that the hiring rate is low. An information strategy that leads to expected consumption falling when labour market prospects improve cannot maximise expected utility when the full information consumption function implies $\frac{dc^*_{mi}}{d\eta} > 0$. The rational inattention aggregate consumption function cannot therefore be downward sloping. This means that if $H'(c_m) < 0$, there will be just one equilibrium.

We therefore need a degree of strategic complementarity for our multiplicity results. This is very

[12]The idea that agents might start with a uniform prior when they first attempt to learn about something, then update, has been suggested as an explanation for experimental decision making results (e.g. Fox and Clemen, 2005). We only differ from that account in that our agents update rationally given an information processing constraint, rather than through a behavioural heuristic.

plausible in models of precautionary saving based on labour market expectations: Ravn and Sterk (2018) show that this complementarity exists as long as labour income risk is countercyclical, which is satisfied if real wages are approximately acyclical. Strategic complementarity often increases the volatility of consumption and labour market variables in response to shocks, but it only leads to multiple equilibria if one or more model equations is sufficiently non-linear. In Ravn and Sterk (2018), the unemployment trap steady state occurs at a hiring rate of zero because the Phillips Curve is kinked at that point by the requirement that vacancy posting cannot be negative. Our model is different because the non-linearity in the consumption function that generates the multiplicity does not come from such an imposed cutoff (though we agree that imposing that vacancies must be positive is sensible), or from some exogenously imposed non-linearity in another part of the model. The non-linearity arises endogenously from optimal household information choices. For this reason, in section 4 we find an unemployment trap with low, but positive, labour market activity, whereas the unemployment trap in Ravn and Sterk has zero employment.

3.8 Dynamic Solution

Our multiplicity result relies on household prior beliefs not being precisely equal to the true equilibrium distribution of the hiring rate. In a model where households process information about the hiring rate each period, they might update their prior beliefs using information processed in previous periods. Our multiple equilibria will only survive in the long run if this prior belief updating does not lead to beliefs collapsing to the true distribution of the hiring rate over time. The model in section 3.1 is static, so there is no room for agents to learn in this way. In this section, we repeat the static problem from that simple model to give households the opportunity to learn about the distribution of the hiring rate. For this section, a 'period' is composed of a morning and afternoon. As in section 3.1, all households are employed each morning, and some lose their job at the end of the morning. Firms then hire new workers for the afternoon in a frictional labour market. As in the static model, households consume all of their income and savings each afternoon, so the income inequality created by the frictional afternoon labour market does not lead to wealth inequality, which would otherwise need to be tracked between periods. Households start from the uniform prior belief in the first period, but they can use information from one period to update their prior beliefs about the distribution of η for the next period. Household prior beliefs are the only link between periods. We show that in this simple model, beliefs do not converge to the true distribution of η, even in the long run. Multiple steady states therefore persist even when households update their prior beliefs over time.

We will consider a simple rule for updating beliefs. If household i reaches the posterior belief $f(\eta_t|c_{mit})$ after processing signals in the morning of period t, we will suppose that their prior belief before information processing in period $(t+1)$ is given by:

$$g_i(\eta_{t+1}|c_{mit}) = \rho f_i(\eta_t|c_{mit}) + (1-\rho), \quad \rho \in [0,1] \tag{16}$$

Intuitively, we take a weighted average of last period's posterior and the uniform $(0,1)$ initial prior, with the weight ρ interpreted as a measure of the (perceived) persistence of the hiring rate η. The particular updating rule is not important, what matters is that households take some of the information processed to arrive at their posterior belief in one period and use it to inform their prior belief in the next period[13].

We will also make one further simplification, which is to assume that when households make their choices in period t they do not take into account the informational value of those choices for future periods. Kreps (1998) introduced this as 'anticipated utility', and Cogley, Colacito and Sargent (2007) showed that in a simple monetary model the solution to this problem is a good approximation to the fully rational model where future information values are taken into account.

With this assumption, the first period optimal $f(\eta_1, c_{mi1})$ (from the uniform prior) is as in figure 1. This is because the anticipated utility assumption means that the agent acts as if they face a series of unconnected static problems, since nothing else connects the periods in this simple example. If the prior is uniform in period 1, the first static problem is exactly the same as in section 3.2. There are two possible posterior beliefs that the households could finish period 1 with, corresponding to the two levels of morning consumption in the optimal menu. If a household chooses the high level of consumption, it is because they have received a signal that the hiring rate is high, and similarly the low value of consumption corresponds to a low signal. There are therefore two possible signal realizations, and so two possible posterior beliefs. Formally, these posteriors are given by the conditional distribution $f(\eta_1|c_{mi1})$. The two posteriors are plotted below.

[13]True Bayesian updating is not possible because that requires that households know the process generating the hiring rate, which is ruled out by our use of ignorance priors.

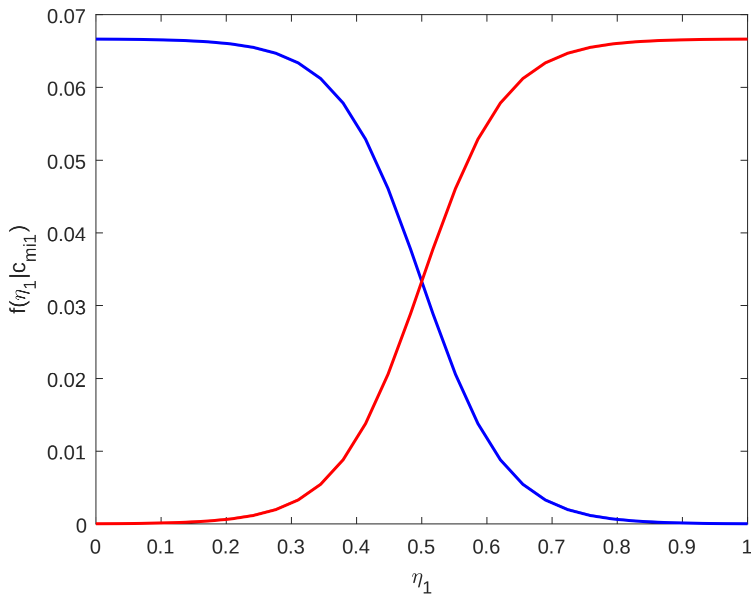

Figure 4: Posterior beliefs from a uniform prior and $\psi = 0.002$ for those who chose low c_m (blue) and high c_m (red)

The period 2 priors will be weighted averages of these posteriors and the initial uniform prior. As long as $\rho > 0$, the households will all start period 2 with a prior with a lower entropy than their period 1 prior. That is, as long as households believe that η has some persistence, they will incorporate some period 1 information into their period 2 prior.

In period 2 this model diverges from the static case. Remember that agents choose how much information to process (κ) such that the marginal benefits of more information equal the marginal cost ψ. Information in prior beliefs is a substitute for information from new signals. This means that the marginal benefits of information are increasing in the entropy of the prior: if prior beliefs are already very informative (low entropy) there is little extra benefit from more information. In period 2, the entropy of each household's prior is lower than it was in period 1. The households therefore process less information in period 2 than they did in period 1. This implies that the households rely more on their priors to guide their decisions in period 2 than they did in period 1, because the period 2 priors contain more information. If they chose a low morning consumption in period 1, their posteriors in period 1 must have suggested that η_1 was low. With little new information processing in period 2, the agent is therefore more likely to choose a low morning consumption in period 2 than a household who chose a higher morning consumption in period 1. This inertia in choices is studied in detail by Matejka, Steiner and Stewart (2017). This is why our model features persistence in individual expectations, as seen in the data.

As long as $0 < \rho < 1$, period 2 priors are more dispersed than period 1 posteriors[14], and are less dispersed than the uniform prior from period 1. The households therefore start period 2 with more precise priors than period 1, but process less information to update from that prior to a posterior. It is therefore not the case that priors degenerate to the true distribution of η over time. In fact, that cannot be the result: each period the households process enough information to return them to the point where the marginal benefit of information equals ψ. We know from period 1 that this occurs before posteriors completely pin down η_t.

The key component of our model that prevents prior beliefs collapsing to the true distribution of the hiring rate is that households don't learn from past realizations of the hiring rate, they learn from past signals they received about η (by assumption they never observe past values of the hiring rate). These signals are costly, and the marginal cost of more information is constant. The marginal benefit of more information falls as priors become more informative, so learning in our model does not lead to beliefs converging to something close to the true distribution of η. If instead households observed the history of realizations of the hiring rate without cost, prior beliefs would converge to the true distribution of the hiring rate[15].

For this simple model, we simulate to find candidate steady states. We choose a level of $\eta_t = \bar{\eta}$ and fix it over time. In period 1, a proportion L choose the low consumption c_l, and $H = (1 - L)$ choose the high c_h, in the optimal menu. The aggregate c_{m1} is simply $Lc_l + Hc_h$. In period 2, that means that a proportion L have priors which bias the agents towards low consumption, and H are biased towards high c. We solve the RI problem for both of these groups, and obtain new posteriors and proportions of households holding each posterior. We iterate this process until aggregate consumption and the composition of beliefs in the population are stable. The graph below plots these candidate steady states when $\rho = 0.9$: the aggregate morning consumption c_m that comes out of this iteration process for each possible $\bar{\eta}$. Again, the equilibrium condition from the firm side of the model 15 is in black, and the model has steady states where the blue set of candidate steady states meets this equilibrium condition.

[14]If $\rho = 1$, households simply take their posterior and use it as the next period's prior without adding any extra noise. In period 1 (with a uniform prior), they process information until the marginal benefit of being informed equals the marginal cost ψ. In the next period, their prior belief would contain the same amount of information as the period 1 posterior. Any further information processing would therefore have a marginal benefit below ψ. If a hiring rate is an equilibrium in the period with a uniform prior, it is therefore trivially a steady state, as no household changes their actions in any future periods. Similarly, if $\rho = 0$ households do not update their prior beliefs between periods and the model is a sequence of unconnected static problems, so any equilibrium in the static model is also a steady state.

[15]To be precise, they would converge to a distribution which is statistically indistinguishable from the true distribution. This is the 'near-rationality' in Adam and Marcet (2011).

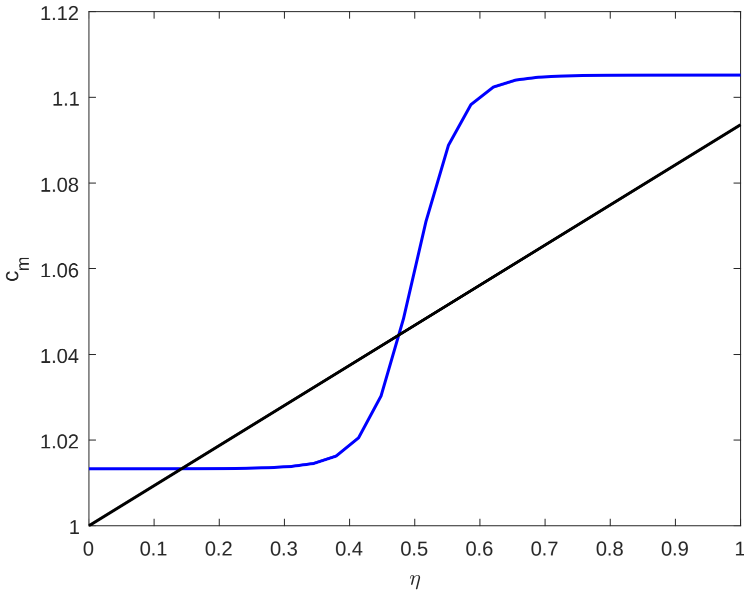

Figure 5: Candidate steady state c_m for each η.

If beliefs did degenerate to the truth, this curve would coincide with the linear full information consumption function[16] in figure 2. In fact we retain the logistic shape of the aggregate response curve from the static problem, so this model generates multiple steady states. There is a high employment steady state with $\eta = 0.48$, and an unemployment trap steady state with $\eta = 0.14$. Note that the blue curve here is closely related to the consumption function in figure 2. The figure 2 consumption function plots the combinations of the hiring rate and aggregate consumption which are consistent with household optimal information and consumption decisions. That is also true for the aggregate consumption function in figure 5, which adds to household optimality that the combination of η_t and c_{mt} must be consistent with a stable distribution of household prior beliefs.

While the distribution of beliefs is stable in steady state, we do not require that every household makes the same choice in every period. There will be churn at the level of individual households, but the distribution of beliefs in the population is constant at every point along the blue steady state aggregate consumption function. The tables below show the transition probabilities between pessimistic (**P**) or optimistic (**O**) prior beliefs at the two steady states.

Table 1: Prior belief transition matrices

(a) Unemployment trap

	$\mathbf{P}_{s+1}$	$\mathbf{O}_{s+1}$
$\mathbf{P}_s$	0.9997	0.0003
$\mathbf{O}_s$	0.9346	0.0654

(b) High employment steady state

	$\mathbf{P}_{s+1}$	$\mathbf{O}_{s+1}$
$\mathbf{P}_s$	0.9612	0.0388
$\mathbf{O}_s$	0.0683	0.9317

[16] We hold the hiring rate constant, so the true distribution of the hiring rate is simply a point. If households know this, with no further information processing they choose the optimum consumption for that hiring rate.

In the high employment steady state, both optimistic and pessimistic beliefs are very persistent, so there is a great deal of belief heterogeneity in steady state. 64% of households hold pessimistic prior beliefs each period, and the remaining 36% hold optimistic beliefs. In the unemployment trap, the hiring rate is so low that even households with optimistic prior beliefs are very likely to receive a signal that η is low, and so to switch to pessimistic beliefs. There is therefore much less heterogeneity in prior beliefs in the unemployment trap steady state: 99.97% of households hold pessimistic beliefs each period.

The amount of belief dispersion in steady state is regulated by the information cost parameter ψ: with a lower information cost households would choose signals with more than two possible realizations, so households would try to distinguish between several different low values of the hiring rate. This would create more belief dispersion in steady states with a low η. This would not remove the multiple steady states: the rational inattention aggregate consumption function would still have the step-like shape displayed in figure 5, there would simply be more steps. In this case rational inattention may generate more than two steady states.

4 HANK model

Here we study the HANK model in Ravn and Sterk (2018) (RS) with the addition of rational inattention to the hiring rate and prior beliefs which do not match the true equilibrium hiring rate distribution. This model is particularly useful because it remains very tractable, despite featuring the uninsurable labour market risk which is necessary to generate a precautionary savings motive.

The tractability is achieved by assuming that there are 'asset-rich' households who are risk neutral, own firms and do not participate in the labour market. The remaining 'asset-poor' households supply labour to firms in a frictional labour market, and cannot borrow. Bonds are in zero net supply, so all asset-poor households hold zero wealth in equilibrium. Employed households are on their Euler equation and are all identical, and unemployed households are hand-to-mouth, consuming their home production. There are therefore only three types of households in the model, not the full distribution seen in other models with uninsurable idiosyncratic risk (e.g. Kaplan, Moll, Violante 2018). Employed asset-poor households are the critical households as they remain on their Euler equation, so they price the bond in equilibrium.

We make one significant change to the Ravn and Sterk model: we remove the constraint that vacancy posting cannot be negative. In Ravn and Sterk's model, this constraint places a kink in the Phillips Curve at a hiring rate of zero, which gives rise to an unemployment trap steady state with zero hiring. Without this, the model has a unique steady state under full information[17]. We show that rational inattention generates multiple steady states in this environment. Other than removing this constraint on vacancies, we make only minor adjustments to the Ravn and Sterk model. Where Ravn and Sterk prevent all households from borrowing, we only apply this borrowing limit to unemployed asset-poor households. With full information, as in Ravn and Sterk's model, this makes no difference to the model as all employed households are identical, so in equilibrium they must still hold zero assets. In contrast, with rational inattention employed households will have heterogeneous beliefs about the future hiring rate, and so will have heterogeneous desires for precautionary saving. This means that employed households will save and borrow small amounts among themselves. There will be a non-degenerate wealth distribution among employed households, but in the calibrated steady state no household accumulates enough wealth to imply that they are on their Euler equation when they become unemployed. All households hit the borrowing constraint when they transition to unemployment, as in the full information model. Bond market equilibrium therefore requires that the net asset position of employed households is zero.

If we maintained the Nash bargaining over wages used by Ravn and Sterk, this wealth distribution would imply that different households would receive different wages, as a wealthier household suffers less from unemployment and so has a more valuable outside option. This is not the focus of this paper, so for simplicity we assume that wages are fixed exogenously.

4.1 Households

Employed households choose how much to consume and save each period based on their current wealth and their expectations about the future hiring rate. Unemployed households face a binding no-borrowing constraint, so consume all of their income from home production each period. Each period an exogenous proportion ω of employed households lose their job, and an endogenous proportion η_s of unemployed households find a job. This is the same as the household problem in Ravn and Sterk, except that here employed households are allowed to borrow[18]. Note that there is no morning/afternoon distinction as in section 3: if a household loses their job at the end of period s, they start to search for new employment in period $s + 1$.

[17]Note that vacancy posting will never be negative in steady state. Removing the explicit constraint simply removes the kink in the Phillips Curve which allows for a zero hiring steady state.

[18]If an employed household chooses to borrow, and then loses their job at the end of that period, they must return to weakly positive saving in their first period of unemployment.

The employed household's problem under perfect information is therefore:

$$\max_{c_{es}} V_s^e = \Big(\frac{c_{es}^{1-\mu}-1}{1-\mu} - \zeta\Big) + \beta \mathbf{E}_s\bigg(\omega(1-\eta_{s+1})V_{s+1}^u\bigg) + \beta \mathbf{E}_s\bigg((1-\omega(1-\eta_{s+1}))V_{s+1}^e\bigg) \quad (17)$$

subject to

$$P_s c_{es} + \frac{b_{s+1}}{R_s} = P_s w + b_s \quad (18)$$

Here V_s^e and V_s^u are the values of being employed and unemployed in period s respectively. Employed households consume c_{es}, and have disutility of labour ζ. The coefficient of risk aversion is μ and the discount factor is β. Prices are P_s, the interest rate is R_s and the exogenous real wage is w. Households save by buying one period bonds b_{s+1}.

The first order condition under perfect information, substituting inflation $\Pi_{s+1} = \frac{P_{s+1}}{P_s}$, is as in RS:

$$c_{es}^{-\mu} = \beta\frac{R_s}{\Pi_{s+1}}\Big(\omega(1-\eta_{s+1})c_{u,s+1}^{-\mu} + (1-\omega(1-\eta_{s+1}))c_{e,s+1}^{-\mu}\Big) \quad (19)$$

Unemployed households are always at their borrowing constraint, and so their problem never matters for equilibrium determination[19].

To get the steady state Euler equation (**EE**) Ravn and Sterk note that $c_{es} = w$ and $c_{us} = \vartheta$ due to the no-borrowing constraint[20], and they substitute for an interest rate Taylor rule:

$$R_s = \max\{\bar{R}\bar{\Pi}^{-\delta_\pi}\Pi^{\delta_\pi}\eta^{\frac{\delta_\eta}{1-\alpha}}, 1\} \quad (20)$$

As they do in drawing their figure 3, we will assume that the interest rate responds only to inflation, that is that $\delta_\eta = 0$.

This gives:

$$1 = \frac{\beta\max\{R^*\Pi^{\delta_\pi}, 1\}}{\Pi}\Big[\omega(1-\eta)\Big(\frac{\vartheta}{w}\Big)^{-\mu} + 1 - \omega(1-\eta)\Big] \quad (21)$$

Where $R^* = \bar{R}\bar{\Pi}^{-\delta_\pi}$.

[19]See RS for a detailed explanation.
[20]ϑ is the payoff to home production.

4.2 Firms

We keep the firm problem very standard. Firms in our model are identical to those in Ravn and Sterk (2018), except they do not face a non-negativity constraint on vacancies. Firms set prices and choose how many vacancies to post each period to maximise profits. There are quadratic price adjustment costs as in Rotemberg (1982), firms are monopolistic and households have CES preferences. The price setting and vacancy posting decisions lead to a Phillips Curve, which in steady state becomes:

$$\phi(1-\beta)(\Pi-1)\Pi = 1-\gamma+\gamma(w+k\eta^{\frac{\alpha}{1-\alpha}}(1-\beta(1-\omega))) \quad (22)$$

Here ϕ measures the extent of price adjustment costs, γ is the elasticity of substitution between goods in the consumer's problem, k is the cost of posting a vacancy, ω is the (fixed) job separation rate, η is the hiring rate and q is the vacancy filling rate, equal to $\eta^{\frac{-\alpha}{1-\alpha}}$, where α is the elasticity of the Cobb-Douglas labour matching function wrt job searchers. $\Lambda_{s,s+1}$ is the discount factor of the owners of the firm, who are assumed to be risk neutral. This steady state Phillips Curve is derived in appendix C.1.

Equation 22 is identical to equation (**PC**) in RS, except that assuming constant wages means that real wage w is not a function of η, and λ_f is dropped. This is the Lagrange multiplier on the constraint that vacancies must be weakly positive. We drop this constraint, so λ_f is always zero in our version of the model. In steady state, vacancies will always be positive in our model.

4.3 Perfect Information Steady State

The steady state Phillips Curve and full information Euler Equation both contain just two endogenous variables: inflation and the hiring rate. These two curves therefore pin down the steady state of the model. They are plotted in figure 6 below, with the Phillips Curve in blue and the Euler Equation for employed households in red[21].

[21] This is drawn using a monthly calibration taken from appendix A3 of Ravn and Sterk (2018). We use this calibration for all of the graphs in this section. The parameter values are detailed in appendix C.2.

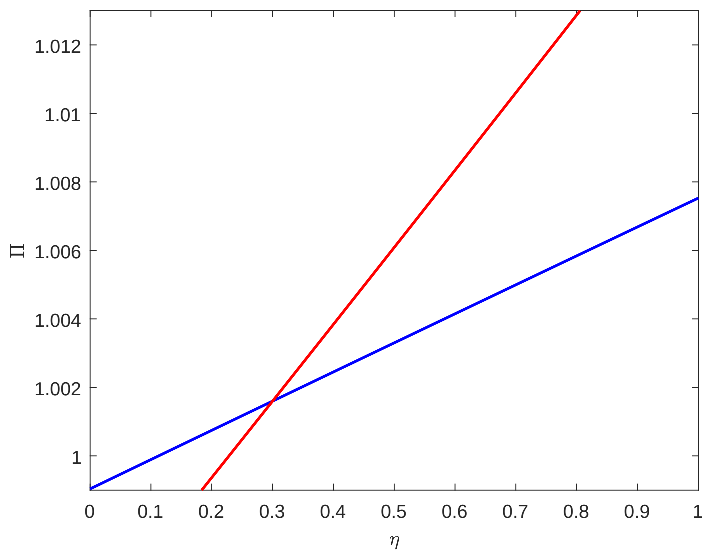

Figure 6: Steady State relations under perfect information. The Phillips Curve is in blue, the steady state Euler equation is in red.

This figure is identical to figure 3 panel 1 in RS, with two exceptions: we have not plotted the steady state relationships at the zero lower bound as we do not study that region in this paper, and we do not have a kink in the Phillips Curve at $\eta = 0$. This is because we have removed the non-negativity constraint on vacancies. In RS the kink implies the existence of a steady state with zero vacancy posting and so zero hiring and employment. With RI we find multiple steady states with $\eta > 0$.

The steady state Phillips Curve is upward sloping because firms face quadratic price adjustment costs (Ascari and Rossi (2012) discuss this result). The Euler Equation is upward sloping because a higher steady state hiring rate η decreases the desire for precautionary saving[22]. To keep the bond market in equilibrium employed households must therefore be encouraged to save more. A higher rate of inflation leads to a higher interest rate, so a higher η is associated with more inflation in steady state.

4.4 Rational Inattention

As in section 3.2, we now amend the household problem so that processing information about the next period hiring rate has marginal cost ψ. Again, we start with uniform prior beliefs $g(\eta) \sim U(0,1)$ for all households. In section 4.5 we allow households to update their prior beliefs using information processed in the previous period, as in section 3.8. The equations of the

[22]To be precise, a higher hiring rate only decreases the precautionary savings motive when labour income risk is countercyclical, which is the case here because we assume that wages are independent of the hiring rate. If labour income risk is procyclical, the Euler equation would slope downwards.

household problem therefore mirror the setup explained in section 3.2. They are set out in detail in appendix C.3.

Relaxing the no-borrowing constraint on employed households means that heterogeneous signals imply a non-degenerate wealth distribution among the employed in steady state. Newly unemployed households may therefore have some savings in some cases, but their Euler equation is still unimportant for the determination of steady state because they are still at their borrowing constraint for all inflation rates and hiring rates consistent with steady state and equilibrium, as in the RS model[23]. If a household is unemployed in period s they will therefore have $b_{s+1} = 0$, regardless of their history prior to period s. Any household moving from unemployment to employment therefore enters the bond market with zero starting wealth.

The solution to the rationally inattentive household problem for an employed household with zero wealth $b_s = 0$ facing a (known) monthly gross inflation of $\Pi = 1.0055$ is plotted as an example in figure 7. It shows how the savings choices of employed households with uniform prior beliefs vary with η_{s+1}.

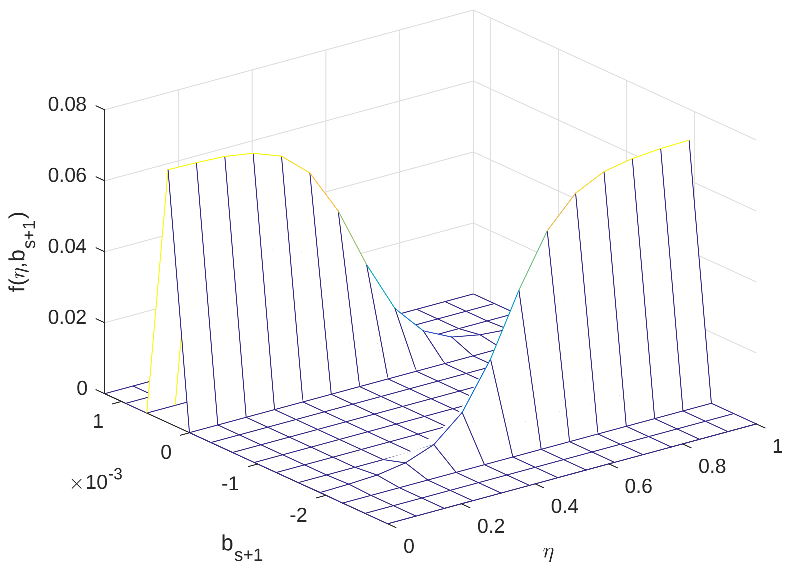

Figure 7: Decision rule for $\Pi = 1.0055$, $b_s = 0$, with a uniform prior belief.

Households choose to limit themselves to a discrete number of savings choices. As in section 3.3, we choose a marginal cost of information such that households restrict themselves to two savings choices, in order to simplify the exposition, but this is not required for our results.

[23]That is, no unemployed household will choose to save if they have any of the saving levels chosen by employed households at any point along the black dashed line in figure 9

As the hiring rate rises, the probability that the household gets a signal that η is high rises. It therefore becomes more likely that they will choose the low level of savings (high consumption) in their optimal menu. This decision rule varies with the household wealth b_s. The more savings a household has in period s, the greater the marginal benefit to saving further (in b_{s+1}) as there are diminishing marginal returns to consumption in each period.

4.5 Belief updating

As in section 3.8 for the static model, we now allow households to update their beliefs over time using the simple rule:

$$g(\eta_{t+2}|b_{i,t+1}) = \rho f(\eta_{t+1}|b_{i,t+1}) + (1-\rho), \quad \rho \in (0,1) \tag{23}$$

That is, household i uses their posterior belief about the hiring rate, which is from information collected in period t (which informed their savings choice $b_{i,t+1}$). They combine this with their original uniform prior, weighting the posterior by ρ and the uniform belief by $1-\rho$. A household with a uniform prior and zero wealth in period t has the decision rule plotted as an example in figure 7, so has two possible realizations of $b_{i,t+1}$. This implies that there are two posteriors they could form in period t, one for each savings choice. This leads to two possible prior beliefs at the start of period $t+1$, which are plotted below. We will refer to a prior belief that results from a signal that η is low in the previous period as *pessimistic* (prior in red below), and the prior belief after a signal that η is high as *optimistic* (in blue).

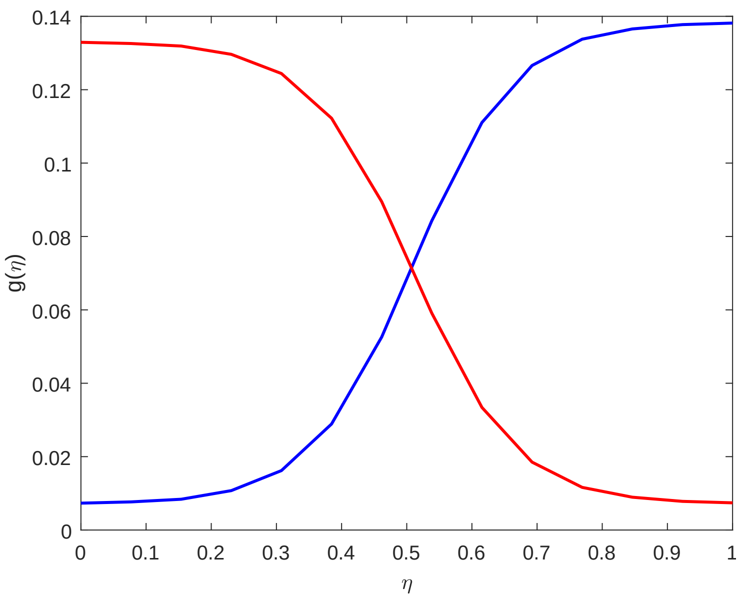

Figure 8: Possible prior beliefs after information processing from a uniform prior.

We solve the household problem using each of these two prior beliefs[24], and find the prior beliefs households could hold in period $t+2$. Iterating this process we find that there are two groups of priors that emerge in steady state, which approximately match the two priors plotted in figure 8: in steady state households are divided into optimists and pessimists[25].

4.6 Rational Inattention Steady States

Compared with the perfect information model in RS, there are two extra requirements for steady state in the rational inattention model: the distributions of wealth and prior beliefs must both be stable. This does not mean that wealth or beliefs are static: each household receives signals with idiosyncratic noise, which implies churn in both wealth and beliefs underneath their stable distributions.

The steady states of the model are plotted below Details of the solution method are given in appendix C.4. The blue and red lines are the Phillips Curve and full information Euler Equation from figure 6. The black dashed line is the steady state Euler Equation with rational inattention, with wealth and belief distributions constant at every point along the line. The calibration is taken from Ravn and Sterk (2018) appendix A3.

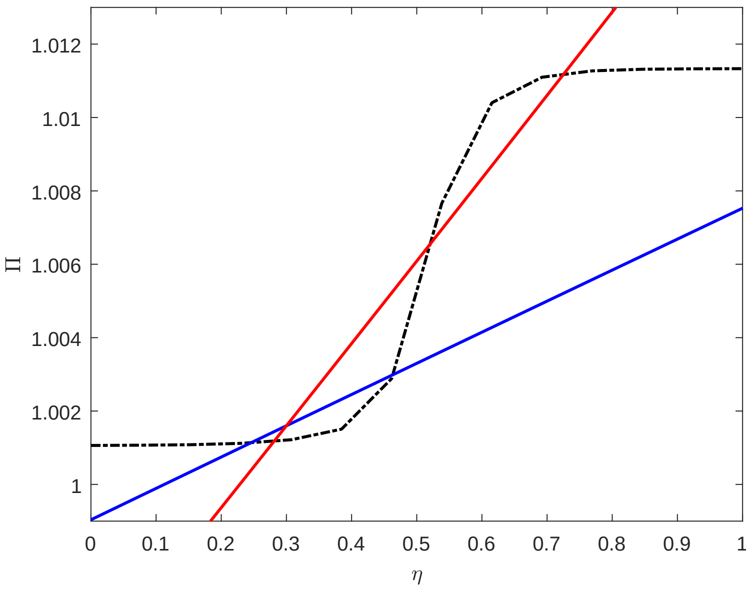

Figure 9: Steady state relations under perfect information and rational inattention ($\psi = 0.0025$)

Under perfect information there is a unique steady state, with $\eta = 0.3$ and $\Pi = 1.00165$. With

[24]The optimal decision rules are plotted in appendix C.4.

[25]We iterate the RI problem until we obtain a set of prior beliefs such that each prior and subsequent information processing implies next-period priors which approximately equal another element of the set. It is in this sense that we obtain two 'groups' of priors rather than two exact priors.

rational inattention there are two steady states, a high employment steady state with $\eta = 0.463$, $\Pi = 1.003$ and a low employment steady state with $\eta = 0.245$, $\Pi = 1.0011$. The annualised inflation rate under full information is 2%. Under rational inattention, the high employment steady state has annual inflation of 3.7%, and the low employment steady state (the unemployment trap) has annual inflation of 1.3%.

A greater hiring rate makes more of the households select the lower savings rate in their menu, and a higher inflation rate encourages more saving through higher interest rates. This is why both the perfect information and rational inattention Euler equations are upward sloping: for net savings to be zero, a higher hiring rate must be offset by higher inflation. The wave shape of the EE curve under rational inattention arises because households do not smoothly adjust their savings in response to the hiring rate, as they do in the perfect information model. At extreme low (or high) η a small change in η does not change household decisions very much, as discussed in section 3.3. Here, this means that the change in inflation required to maintain equilibrium and steady state is small, and so the EE curve is flat for very low (high) hiring rates. The reverse is true for η in the middle of its range: savings respond more to η in that range under rational inattention than under perfect information, and so the EE curve is steep there. The logistic shape of the probabilities in the optimal decision rule (see figure 7) therefore drives the multiplicity.

As in the static model of section 3, the flat portions of the steady state Euler equation at extreme low or high values of the hiring rate feature less belief dispersion than values of η in the middle of its support. Very few households receive signals that the hiring rate is high when it is in fact below 0.3, and so very few households have optimistic prior beliefs in steady state. As in section 3.8, the amount of belief dispersion in steady state is regulated by the information cost parameter ψ: with a lower information cost households would choose signals with more than two possible realizations, so households would try to distinguish between several different low values of the hiring rate. This would create more belief dispersion in steady states with a low η, and would possibly generate more than two steady states.

As discussed in section 4.5, in steady state with our particular choice of information cost ψ priors fall into two groups of similar beliefs. The graph below plots those two prior beliefs present in steady state. For this information processing cost, this is very similar to the two prior beliefs seen in the period after information processing from a uniform prior, plotted in figure 8. That is, a household with a pessimistic prior belief is very likely to receive a signal that η is low, and that will be such that their prior belief does not adjust significantly from its initial position.

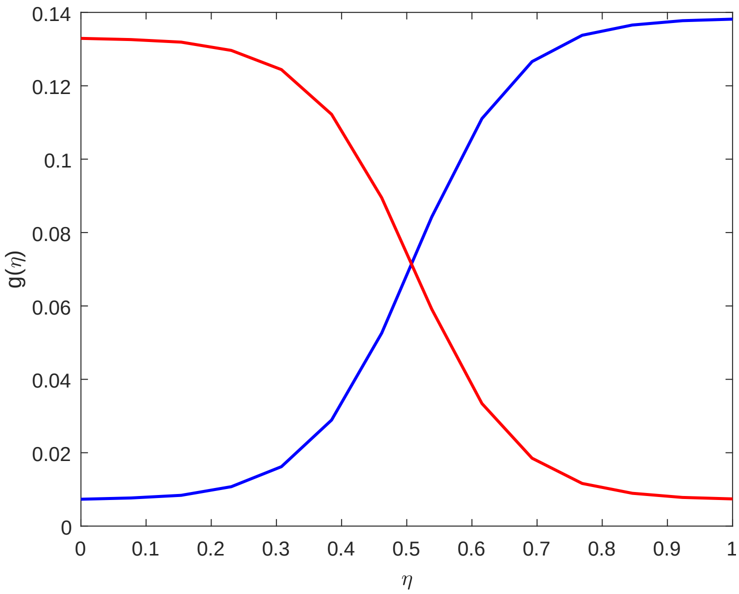

Figure 10: Steady state relations under perfect information and rational inattention ($\psi = 0.0025$)

There is churn underlying the steady states in figure 9. The tables below give the transition probabilities between the two prior beliefs at the unemployment trap and high employment steady states respectively:

Table 2: Prior belief transition matrices

(a) Unemployment trap

	$\mathbf{P}_{s+1}$	$\mathbf{O}_{s+1}$
$\mathbf{P}_s$	0.996	0.004
$\mathbf{O}_s$	0.840	0.160

(b) High employment steady state

	$\mathbf{P}_{s+1}$	$\mathbf{O}_{s+1}$
$\mathbf{P}_s$	0.968	0.032
$\mathbf{O}_s$	0.208	0.792

At both steady states, a household with pessimistic prior beliefs ($\mathbf{P}$) is very likely to receive a signal that the hiring rate is low, and so to remain pessimistic. In the high employment steady state, an optimistic household ($\mathbf{O}$) is also very likely to remain optimistic. Hiring rate expectations are therefore heterogeneous and persistent, as in the data. In the unemployment trap steady state, the hiring rate is very low, so even households with optimistic prior beliefs are very likely to receive signals that the hiring rate is low. The unemployment trap therefore features a very high proportion (99.5%) of households on pessimistic prior beliefs, while the high employment steady state has much greater dispersion of beliefs: 86.7% of households hold pessimistic beliefs, and the remaining 13.3% are optimistic.

5 Survey Expectations

In the Survey of Consumer Expectations, employed households are asked the following question:

> Suppose you were to lose your main job this month. What do you think is the percent chance that within the following 3 months, you will find a job that you will accept, considering the pay and type of work?

This is precisely the hiring rate studied in our models, and in the model of Ravn and Sterk (2018). In this section we show that the survey responses display several features which are present in our model. The figure below shows the histogram of responses as deviations from the mean response that month.

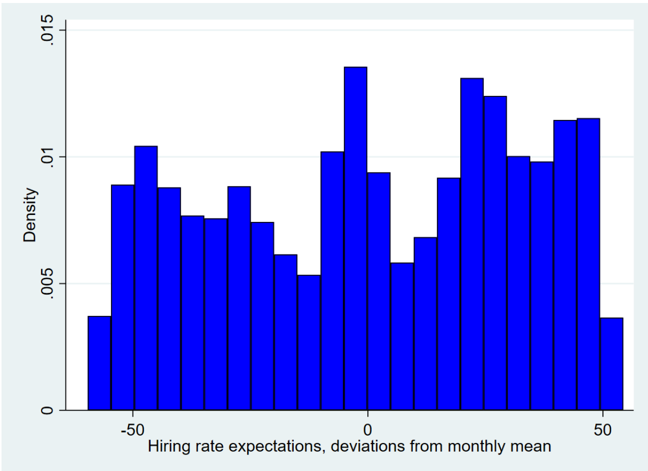

Figure 11: Histogram of hiring rate expectations from the Survey of Consumer Expectations, deviations from the average for the month of the interview.

There is a great deal of dispersion. In the average month the mean response is 53.5%, and the standard deviation is 32.1%. However, this disagreement could all be due to household heterogeneity, not differences in information. High and low education households, for example, could agree exactly about aggregate labour market conditions, but expect different hiring rates because they are making predictions for different segments of that labour market. To explore this, we run the following regression:

$$\mathbf{E}_{it}\eta_{it,t+3} = \alpha_0 + \alpha_1 X_{it} + \varepsilon_{it} \tag{24}$$

The dependent variable is household i's expectation of their own hiring rate for the months from t to $t+3$. X_{it} collects every personal characteristic available in the SCE[26]. The R^2 for this regression is 0.084. Adding time fixed effects to pick up any variation due to macroeconomic factors

[26] The controls are age, age^2, income, income^2, education, gender, race, job tenure, financial distress, state, home ownership and number of times the household has been in the survey. Financial distress is measured as the percentage chance that the household will struggle to pay their bills in the next three months.

only increases this a small amount, to 0.094. The vast majority of heterogeneity in labour market expectations does not come from the dimensions of household heterogeneity recorded in the SCE. There are of course other dimensions of household heterogeneity which are not collected in the survey, which could explain more of the heterogeneity, but it is unlikely that this would account for all of the currently unexplained heterogeneity.

Beliefs could, however, be heterogeneous and still co-ordinate on a high (or low) unemployment steady state. If all households agree in some period that the future hiring rate has fallen by 5% relative to their expectations in the previous period, that would give the same self-fulfilling expectations mechanism described in Ravn and Sterk (2018) and Heathcote and Perri (2018), even if the cross-sectional distribution of beliefs features lots of heterogeneity each period. To explore this, we use the panel nature of the SCE to study changes in the unexplained part of expectations, $\Delta\varepsilon_{it}$. If there is strong co-ordination in beliefs, then changes in household-level expectations should be largely explained by changes in average expectations. Consider the following model:

$$\Delta\varepsilon_{it} = \beta_0 + \beta_1 \Delta\bar{\varepsilon}_t + u_{it} \tag{25}$$

Here $\bar{\varepsilon}_t$ is the average unexplained hiring rate expectation across households in period t. By taking expectations over households i we obtain:

$$\Delta\bar{\varepsilon}_t = \beta_0 + \beta_1 \Delta\bar{\varepsilon}_t + \bar{u}_t \tag{26}$$

By definition, therefore, $\beta_0 = 0$ and $\beta_1 = 1$. If households do agree on movements in the hiring rate but have different (stable) degrees of optimism or pessimism, the first differencing in this model would remove the individual-specific constants, and we would see a very high R^2. Fitting this model to the survey data, we in fact obtain an R^2 of just 0.0017. Less than 0.2% of the variation in revisions to household expectations is explained by movements in average expectations. There is therefore a great deal of idiosyncratic variation in *updates* to hiring rate expectations, and only weak co-ordination of beliefs. This is a natural feature of our model, because households form their expectations by observing very noisy signals.

Our model also produces a large amount of persistence in household beliefs through the updating of prior beliefs. This can be seen most clearly in table 2b above: even in the high employment steady state, which features dispersed expectations, individual households are extremely unlikely to revise their beliefs from one period to the next. To investigate this persistence in the data, we run the following regression:

$$\varepsilon_{it} = \gamma_1 \varepsilon_{it-1} + W_t + e_{it} \tag{27}$$

Here W_t are month fixed effects, which pick up any variation in expectations due to macroeconomic factors, including average beliefs. The results are displayed in table 4 below.

Table 3: The unexplained part of hiring rate expectations is persistent at the household level.

	(1) Hiring Rate Residual
L.Hiring Rate Residual	0.719***
	(0.00390)
Observations	31830
R^2	0.521

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The coefficient γ_1 is large and significant. This implies that hiring rate expectations are indeed persistent at the household level, as predicted by our model.

A final observation from this data is that the variance of expectations is an order of magnitude larger than the variance of the actual (aggregate) hiring rate, calculated using labour flows in the CPS. Estimating an AR(1) model on the hiring rate, we estimate the long run standard deviation[27] of this series at 4.7%. In contrast, the standard deviation of the unexplained part of hiring rate expectations ϵ_{it} is 30.7%. If households knew the true distribution of the hiring rate and received noisy signals (as in a standard rational inattention model a la Sims (2003)), we would expect the standard deviation of beliefs to be *lower* than the standard deviation of the data[28]. The data therefore suggests that households do not know the true equilibrium distribution of the hiring rate, and instead hold prior beliefs which are significantly more dispersed than that true distribution[29]. This supports our assumption that households hold 'ignorance priors' about the future hiring rate.

[27]The hiring rate is found to be stationary at the 1% level. Labour flows data begins in 1990.

[28]This follows from the Law of Total Variance. If household i collects a signal s_i about the hiring rate η, they form a posterior expectation $\mathbf{E}(\eta|s_i)$. The variance of these conditional expectations across i is given by $V(\mathbf{E}(\eta|s_i)) = V(\eta) - \mathbf{E}(V(\eta|s_i))$. If prior beliefs are correct, the unconditional variance $V(\eta)$ is the true variance of the hiring rate, and so the variance of expectations must be weakly less than this true variance.

[29]This statement relies on our conditioning variables picking up the majority of expectation variation due to heterogeneity in the actual hiring rates experienced by households. Since the variance of expectations is an order of magnitude larger than that of the observed average hiring rate, we are confident that household beliefs are more dispersed than the true distribution of the labour market variables they are trying to estimate.

6 Conclusion

We have proposed a model of a self-fulfilling expectation driven unemployment trap which is driven by rational inattention. The model features households who face costs of processing information about the future hiring rate, and who do not know the true equilibrium distribution of that rate. The households optimally choose to process noisy signals which imply a highly nonlinear response of aggregate consumption to changes in labour market conditions, which leads to the possibility of multiple steady states. We have shown that in the HANK model of Ravn and Sterk (2018) there is a unique steady state with positive hiring if households have full information. Rational inattention generates two steady states: a high employment steady state and an unemployment trap with low, but positive, inflation and hiring activity. Expectations in the model are consistent with key properties of survey expectations: labour market beliefs are heterogeneous, persistent, and display greater variance than that implied by accurate prior beliefs.

References

[1] Adam, Klaus ; Marcet, Albert (2011) *Internal rationality, imperfect market knowledge and asset prices* Journal of Economic Theory, 2011, Vol.146(3), pp.1224-1252

[2] Adam, Klaus ; Kuang, Pei ; Marcet, Albert (2012) *House Price Booms and the Current Account.* NBER Macroeconomics Annual 2012, Vol.26(1), pp.77-122

[3] Andrade, Philippe ; Gaballo, Gaetano ; Mengus, Eric ; Mojon, Benoît (2019) *Forward guidance and heterogeneous beliefs.* American Economic Journal: Macroeconomics, forthcoming

[4] Armantier, Olivier ; Bruine de Bruin, Wändi ; Topa, Giorgio ; Klaauw, Wilbert ; Zafar, Basit (2015) *Inflation Expectations and Behavior: Do Survey Respondents Act on Their Beliefs?* International Economic Review, Vol. 56, No. 2, May 2015

[5] Ascari, Guido ; Rossi, Lorenza (2012) *Trend Inflation and Firms Price-Setting: Rotemberg Versus Calvo.* Economic Journal, September 2012, Vol.122(563), pp.1115-1141

[6] Beaudry, Paul ; Galizia, Dana ; Portier, Franck (2017) *Reconciling Hayek's and Keynes' Views of Recessions.* The Review of Economic Studies, 2018, Vol. 85(1), pp.119-156

[7] Challe, Edouard ; Ragot, Xavier (2016) *Precautionary Saving over the Business Cycle.* The Economic Journal, 126 (February), 135–164

[8] Challe, Edouard ; Matheron, Julien ; Ragot, Xavier ; Rubio-Ramirez, Juan Francisco (2017) *Precautionary Saving and Aggregate Demand.* Quantitative Economics 8 (2017), 435–478

[9] Cogley, Timothy ; Colacito, Riccardo ; Sargent, Thomas (2007) *Benefits from U.S. Monetary Policy Experimentation in the Days of Samuelson and Solow and Lucas* Journal of Money, Credit, and Banking, Feb 2007, Vol.39(s1), pp.67-991

[10] Cooper, Russell ; John, Andrew (1988) *Coordinating Coordination Failures in Keynesian Models.* The Quarterly Journal of Economics, Vol. 103, No. 3 (Aug., 1988), pp. 441-463

[11] Diamond, Peter A. (1982) *Aggregate Demand Management in Search Equilibrium.* Journal of Political Economy, 1982, vol. 90, no. 5

[12] Fox, Craig R. ; Clemen, Robert T. (2005) *Subjective Probability Assessment in Decision Analysis: Partition Dependence and Bias toward the Ignorance Prior* Management Science, 1 September 2005, Vol.51(9), pp.1417-1432

[13] Heathcote, J., and Perri, F. (2018) *Wealth and Volatility.* The Review of Economic Studies, 2018, Vol. 85(4), pp.2173-2213

[14] Jung, Junehyuk ; Kim, Jeong-Ho ; Matějka, Filip ; Sims, Christopher A (2015) *Discrete Actions in Information-constrained Decision Problems.* Working Paper

[15] Kreps, David (1998) *Anticipated Utility and Dynamic Choice* 1997 Schwartz Lecture, in "Frontiers of Research in Economic Theory", Edited by Donald P. Jacobs, Ehud Kalai, and Morton Kamien, Cambridge University Press, Cambridge, England.

[16] Luo, Yulei (2008) *Consumption Dynamics under Information Processing Constraints.* Review of Economic Dynamics, 2008, Vol.11(2), pp.366-385

[17] Luo, Yulei; Nie, Jun; Wang, Gaowang; Young, Eric R. (2017) *Rational Inattention and Dynamics of Consumption and Wealth in General Equilibrium* Journal of Economic Theory, November 2017, Vol.172, pp.55-87

[18] Maćkowiak, Bartosz ; Wiederholt, Mirko (2009) *Optimal Sticky Prices under Rational Inattention* American Economic Review, 99, 769-803

[19] Maćkowiak, Bartosz ; Wiederholt, Mirko (2015) *Business Cycle Dynamics under Rational Inattention* Review of Economic Studies, 82, 1502-1532.

[20] Matějka, Filip (2017) *Rationally Inattentive Seller: Sales and Discrete Pricing.* Review Of Economic Studies Volume 83 Issue 3 Page 1125

[21] Matějka, Filip, McKay, Alisdair (2015) *Rational inattention to discrete choices: A new foundation for the multinomial logit model.* American Economic Review, 1 January 2015, Vol.105(1), pp.272-298

[22] Matějka, Filip ; Steiner, Jakub ; Stewart, Colin (2017) *Rational Inattention Dynamics: Inertia and Delay in Decision Making* Econometrica, March 2017, Vol.85(2), pp.521-553

[23] Meeks, Roland ; Monti, Francesca (2018) *Heterogeneous beliefs in the Phillips curve*, Working Paper

[24] Morris, Stephen ; Shin, Hyun Song (2000) *Rethinking multiple equilibria in macroeconomic modeling.* NBER Macroeconomics Annual, 2000, Vol.15(1), pp.118-161

[25] Ravn, Morten O. ; Sterk, Vincent (2018) *Macroeconomic Fluctuations with HANK & SAM: An Analytical Approach.* Working Paper

[26] Rotemberg, Julio J. (1982) *Sticky Prices in the United States.* Journal of Political Economy, 1 December 1982, Vol.90(6), pp.1187-1211

[27] Sims, Christopher A. (2003) *Implications of Rational Inattention.* Journal of Monetary Economics, 2003, Vol.50(3), pp.665-690

[28] Sims, Christopher A. (2006) *Rational inattention: Beyond the linear-quadratic case.* American Economic Review, 2006, Vol.96(2), pp.158-163

[29] Wiederholt, Mirko (2017) *Empirical Properties of Inflation Expectations and the Zero Lower Bound*, Working Paper

A Aggregate consumption function graphs for the static model

The graph below is the same as figure 2, with an extra curve added in black. This is the aggregate consumption function with a lower (but still positive) cost of information $\psi = 0.00064$. This implies $\kappa \approx 1$.

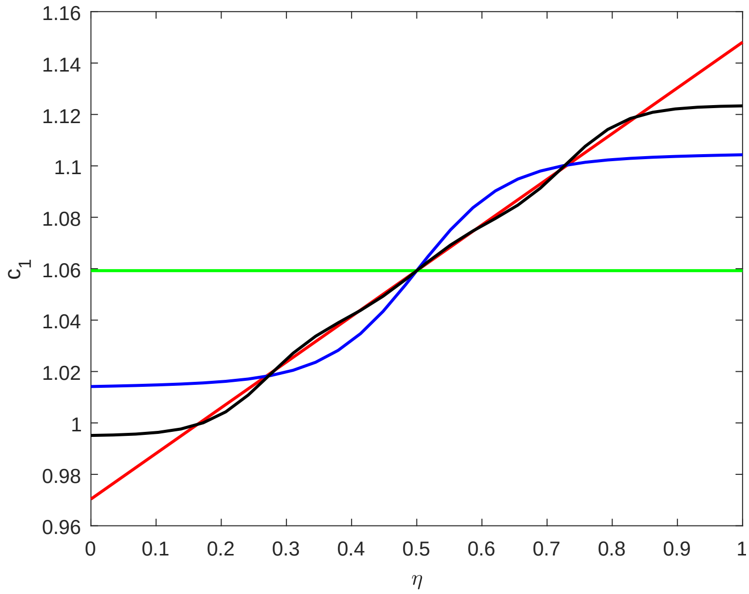

Figure 12: Aggregate consumption function with $\kappa = 0$ (green), $\kappa = 0.5$ (blue), $\kappa = 1$ (black) and in the unconstrained case (red)

The shape of the aggregate response curve in the less constrained ($\psi = 0.00064$) case has the same form as the baseline case of $\psi = 0.002$, but with this greater information processing capacity agents choose from four levels of consumption, so there are four flat regions in the aggregate response curve.

B Static Model: Firms

Here we derive the linear equilibrium relationship between aggregate consumption and the hiring rate in section 3.1.

In period 1, all households are employed at wage 1 and the firm sells $\bar{c}_1$ units of the consumption good. Assume that any unsold output in period 1 is wasted. Firm profits are therefore equal to $\bar{c}_1 - 1$.

New hires in period 2 (h) are determined by the number of vacancies posted (v) and the number of job seekers (u) through the matching function:

$$h = mv^{1-\alpha}u^{\alpha} \tag{28}$$

Now note that the number of job seekers at the start of period 2 is equal to the number of separations at the end of period 1, ω, since all households were employed in period 1. In order to hire, firms must post vacancies. The cost per vacancy is k. Using the matching function, we have that the cost of hiring one worker is[30]:

$$C(\eta) = km^{\frac{-1}{1-\alpha}}\omega\eta^{\frac{1}{1-\alpha}} \tag{29}$$

Assume that these costs must be paid out of period 1 profits, before firms do any period 2 production or sales. The profit per hire is:

$$D(\eta) = F - w - C(\eta) \tag{30}$$

Where F is the value of production from that worker, equal to period 2 aggregate consumption per worker plus the value of inventory at the end of period 2. Assume that this is sufficiently high that $D(\eta) > 0$ for all $\eta \in (0,1)$. That is, in period 2, workers are very productive, and any output not sold is held as inventory, which has a high value to the firm. The firm would therefore always like to hire as many workers as possible, given the working capital constraint that they must use period 1 profits to pay for vacancies.

This creates an upward sloping relationship between $\bar{c}_1$ and η: when period 1 consumption is higher, the firm sells more in period 1, and makes more profit. That means the firm can post more vacancies in period 2, and so the hiring rate increases. Specifically, the hiring rate is pinned down by:

$$\bar{c}_1 - 1 = km^{\frac{-1}{1-\alpha}}\omega\eta^{\frac{1}{1-\alpha}} \tag{31}$$

In the graphs in section 3.1 I further assume that $\alpha = 0$, so the matching function only depends on the number of vacancies posted, which implies a linear relationship between $\bar{c}_1$ and η:

$$\bar{c}_1 = 1 + \frac{k\omega}{m}\eta \tag{32}$$

This is the linear relationship used in section 3.1. The parameters used there are $\beta = 0.975$, $R = \frac{1}{\beta}$, $\omega = 0.4$, $m = 0.25$, $k = 0.045w$, $w = 1.3$, $\theta = 0.4$

[30] I am assuming that firms are large, so the proportion of vacancies filled equals the probability of filling a vacancy.

C HANK model details

C.1 Firm problem

Firm j in period s chooses price P_{js} and vacancy posting v_{js} to maximise real profits. The firm's costs are wages, vacancy posting costs, and quadratic price adjustment costs:

$$\max_{P_{js},v_{js}} \mathbf{E}_t \sum_{s=t}^{\infty} \Lambda_{j,t,t+s} \left[\frac{P_{js}}{P_s} y_{js} - w_s n_{js} - k v_{js} - \frac{\phi}{2} \left(\frac{P_{js} - P_{j,s-1}}{P_{j,s-1}} \right)^2 \right] \tag{33}$$

subject to

$$y_{js} = \left(\frac{P_{js}}{P_s} \right)^{-\gamma} y_s \tag{34}$$

$$y_{js} = \exp(A_s) n_{js} \tag{35}$$

$$n_{js} = (1-\omega) n_{j,s-1} + v_{js} q_s \tag{36}$$

Here y_{js} is firm j output and $y_s = \int y_{js} dj$ is total output. Firm j employs n_{js} units of labour. Each period it loses a fraction ω of its workforce through exogenous separations. The vacancy filling rate is q_s, so $v_{js} q_s$ is the number of new hires by firm j in period s.

Solving this profit maximisation we have:

$$1 - \gamma + \gamma mc_s = \phi(\Pi_s - 1)\Pi_s - \phi \mathbf{E}_s \Lambda_{s,s+1} \left[\frac{y_{s+1}}{y_s} (\Pi_{s+1} - 1) \Pi_{s+1} \right] \tag{37}$$

where

$$mc_s = \frac{w_s}{e^{A_s}} + \frac{k}{q_s} - (1-\omega) \mathbf{E}_s \Lambda_{s,s+1} \frac{k}{q_{s+1}} \tag{38}$$

Here we have dropped the j subscripts because all firms make the same choices.

The labour market matching function is:

$$M(u_s, v_s) = u_s^{\alpha} v_s^{1-\alpha} \tag{39}$$

The vacancy filling rate q_s and hiring rate η_s are given by $\frac{M_s}{v_s}$ and $\frac{M_s}{u_s}$ respectively, so:

$$q_s = \eta_s^{\frac{-\alpha}{1-\alpha}} \tag{40}$$

Also note that the firm owners are risk neutral, so the stochastic discount factor $\Lambda_{s,s+1} = \beta$ for all periods s. The technology shock A_s equals 0 in steady state. The steady state Phillips Curve is therefore:

$$\phi(1-\beta)(\Pi-1)\Pi = 1-\gamma+\gamma(w+k\eta^{\frac{\alpha}{1-\alpha}}(1-\beta(1-\omega))) \tag{41}$$

C.2 Calibration

These parameter values are used to draw the figures in section 4. The calibration is monthly, and is taken from Ravn and Sterk (2018) section A3. The only parameter not taken from their work is ψ, the marginal cost of information. This is set to ensure that each household always chooses an optimal menu of savings choices with two discrete savings levels.

Parameter	**Name**	**Value**
ε	Elasticity of substitution	6
ϕ	Price adjustment cost	96.674
k	Vacancy posting cost	$0.05w$
α	Elasticity of matching function	0.5
ω	Monthly separation rate	0.02
w	Wage	0.8332
θ	Home production value	$0.8w$
β	Discount factor	0.9914
μ	Coefficient of risk aversion	2
ψ	Information processing cost	0.0025

C.3 Household problem with rational inattention

We substitute the budget constraint 18 into the value function, and add in information costs. The employed household problem with RI is therefore as follows:

$$\begin{aligned} f &= \arg\max_{\hat{f}} \mathbf{E}[V_s^e(\eta_{s+1}, b_{s+1})] - \psi\kappa \\ &= \arg\max_{\hat{f}} \int_{\eta_{s+1}} \int_{b_{s+1}} V_s^e(\eta_{s+1}, b_{s+1}) \hat{f}(\eta_{s+1}, b_{s+1}) d\eta_{s+1} db_{s+1} - \psi\kappa \end{aligned} \tag{42}$$

subject to

$$\int_x \hat{f}(\eta_{s+1}, b_{s+1})dx = g(\eta_{s+1}) \qquad \forall \eta_{s+1} \tag{43}$$

$$\hat{f}(\eta_{s+1}, b_{s+1}) \geq 0 \qquad \forall \eta_{s+1}, b_{s+1} \tag{44}$$

$$H[g(\eta_{s+1})] - \mathbf{E}_{b_{s+1}} H[\hat{f}(\eta_{s+1}|b_{s+1})] \leq \kappa \tag{45}$$

The function $H[.]$ is the entropy of the distribution over which it operates. It is defined in equation 13. These constraints are explained in section 3.2. As in the static model, we will consider values for the information cost ψ that imply households optimally limit themselves to a menu with two choices of saving b_{s+1}. This makes the diagrams clearer, but it is not necessary for our results. A lower information cost would imply more choices in the optimal menu, and potentially more steady states than the two we find in this section.

We assume that agents know the current value of all parameters and variables except the hiring rate. They can observe current interest rates and prices. Here we consider the problem in steady state, so we assume agents make their decisions expecting inflation to remain constant. This implies (through the interest rate rule in equation 20) a constant interest rate. Households either know this or simply expect interest rates to remain constant as they do with inflation. We normalise current prices P_s to 1, so current inflation enters the problem through R_s and assuming constant inflation pins down expectations of P_{s+1}. We run this problem many times, for a variety of possible inflation rates.

Specifically, households solving this problem are processing information about the hiring rate in the next period. For simplicity, we do not allow them to process information about the hiring rate in periods further in the future.

This implies that unemployed households derive no benefit from information about the hiring rate, since in the current period they are at their borrowing constraint. Information about the next period hiring rate has no effect on their decision. This means that unemployed households choose to process no information at all. In steady state the prior belief distribution of newly hired households is therefore identical to that of the population as a whole.

C.4 Finding steady states with belief updating

The graphs below plot the decision rules of a household with $b_s = 0$ in two cases. The first plot is for a household with the pessimistic prior belief from figure X, and the second is for a household with the optimistic prior belief from figure X.

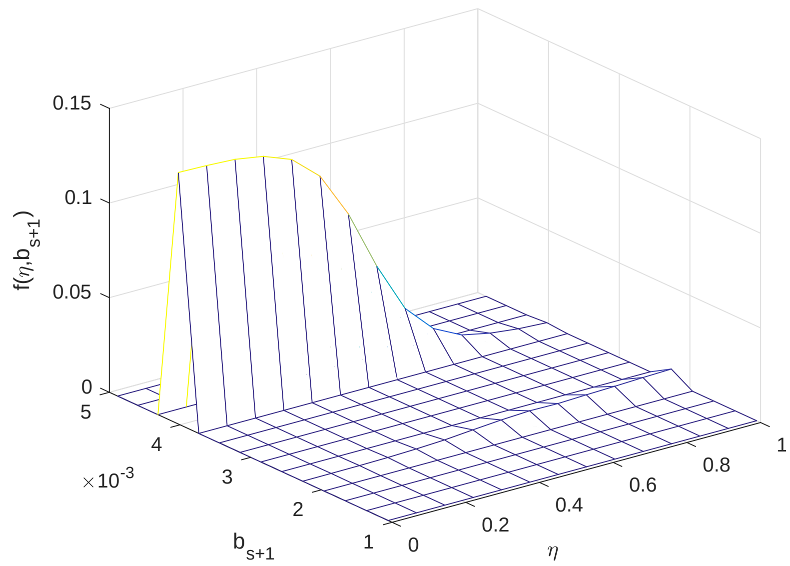

Figure 13: Decision rule for $\Pi = 1.0055$, $b_s = 0$, with a pessimistic prior belief.

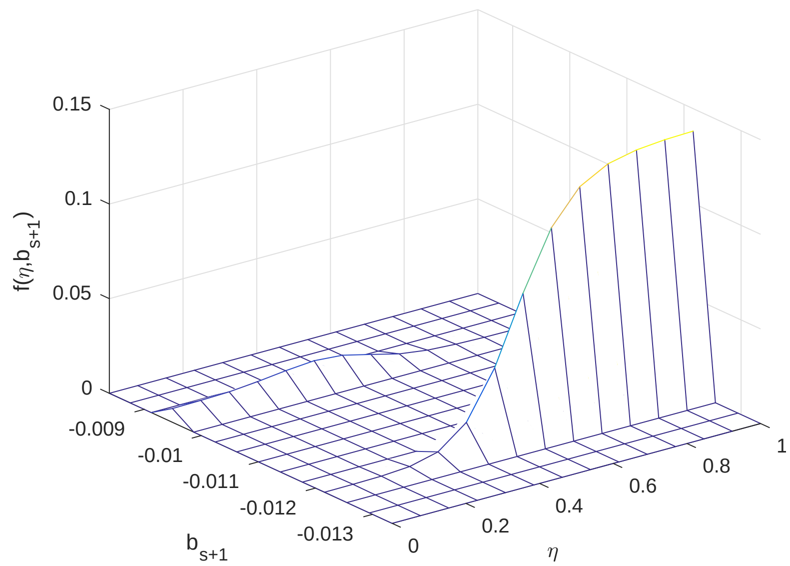

Figure 14: Decision rule for $\Pi = 1.0055$, $b_s = 0$, with an optimistic prior belief.

To make progress in solving for the stable wealth and belief distributions, we make an approximation. Consider the decision rules (solution to the RI problem) plotted below. They are for two households with optimistic prior beliefs and different levels of wealth. The first decision rule is for $b_s = 0$, and the second is for $b_s = 0.014$, which is the highest wealth seen in the steady states we compute. The scale on the optimal savings choice (b_{s+1}) axis has been normalised so that full information optimal savings lie between 0 and 1 in both graphs.

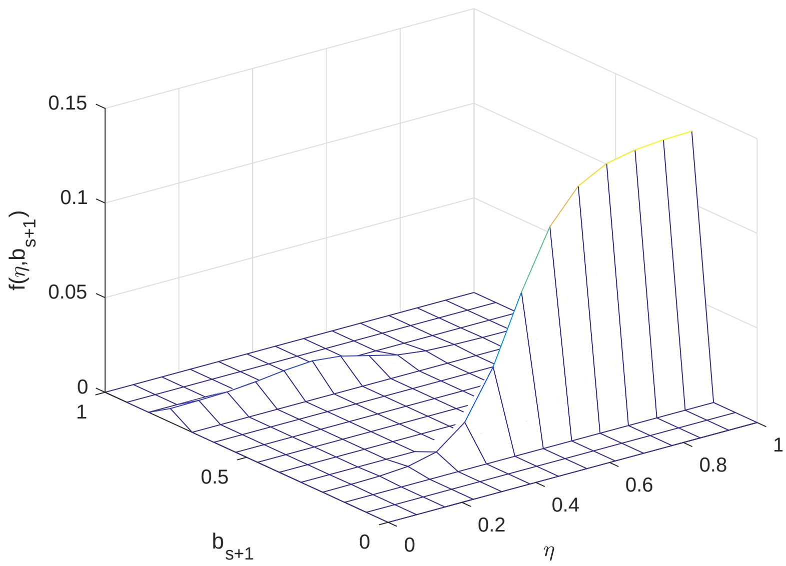

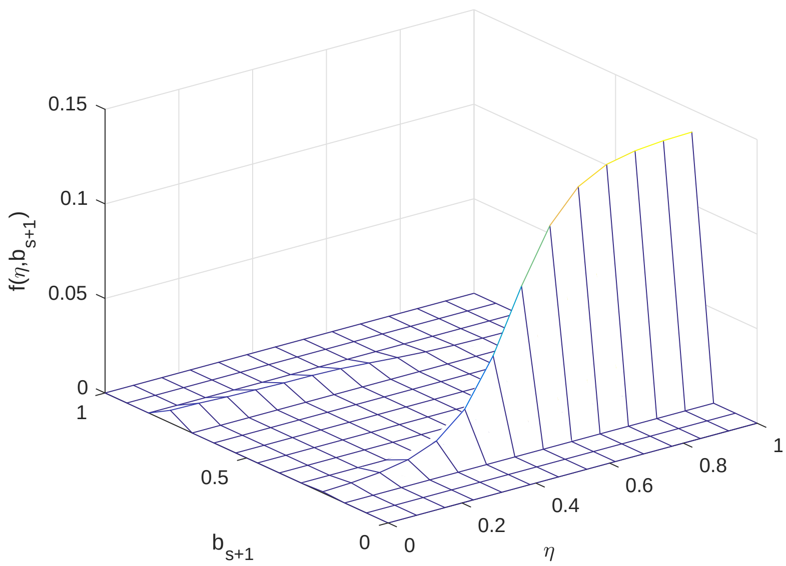

Without normalisation these two decision rules would have different scales on the b_{s+1} axis. The richer household chooses to save more than the poorer household in all states of the world. With the normalisation, we can see that the two decision rules look very similar. That is, the *information strategies* of the two households are very similar: the two possible posterior beliefs about η in each household's optimal menus are almost identical. The approximation we make is that these information strategies are in fact identical. The *action strategies*, how the posteriors are mapped into choices of b_{s+1}, will differ with wealth. Intuitively, we think of a household contracting their information processing to some outside agent. The information processor is given the household's prior beliefs, and the marginal cost of information, but not the household's wealth. The processor receives a signal and forms a posterior $\hat{\eta}$, which it reports back to the household. The household then decides how much to save given that posterior estimate. In this way wealth does not enter the information strategy (the processor's signal collection) but it does enter the action strategy (the household savings choice).

This approximation is small, but the steady states of this approximated model are significantly easier to find than the steady states of the un-approximated model. With this approximation, the posterior beliefs about η implied by the two RI solutions are the same, even though they imply different savings choices for the two households. We can therefore find the steady state distribution of prior beliefs for the population by taking a household with $b_s = 0$, and iterating the belief updating procedure until the steady state belief distribution is constant. We use $b_s = 0$ because, by bond market clearing, this will be the average wealth of employed households in steady state. After finding this steady state belief distribution for a particular combination of inflation and the hiring rate, we use it to find the steady state wealth distribution. This is significantly easier than trying to jointly determine the two distributions. For each inflation rate, the net asset position of the population of employed households is monotonically decreasing in the hiring rate as the precautionary savings motive weakens. There is therefore a unique hiring rate for each level of inflation which is consistent with a net asset position of zero, as required for bond market clearing. The set of these Π, η pairs is the steady state Euler equation under rational inattention.

The approximation mentioned above does not mean that wealth and information are independent: households who receive a signal that η is low will save a lot, and so will become wealthier. They will have prior beliefs in future periods which are biased towards low η. There is therefore feedback from signals and prior beliefs to wealth. The mechanism that our approximation removes is that wealth, in turn, affects the optimal amount of information processed, and so affects beliefs in future periods. This link from wealth to information is small, as shown in figures **??** and **??**.