Rational Inattention and the Unemployment Trap*

Martin Ellison[†] and Alistair Macaulay[‡]

February 13, 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 the Great Recession. Existing models of this mechanism have households co-ordinating their labour market expectations on a particular equilibrium, as is common in models with multiple equilibria (Morris and Shin, 2000).

*This paper supersedes an earlier paper titled 'Rational Inattention and Multiple Equilibria'. We thank Paul Beaudry, Gianluca Femminis, Filip Matějka, Albert Marcet, Edoardo Palombo, Morten Ravn, Ricardo Reis, Christopher Sims, and Mirko Wiederholt for their 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

In this paper, we start by showing that such belief co-ordination is unrealistic. In the Survey of Consumer Expectations (SCE), 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 hiring rate expectation is crucial for precautionary savings in models with frictional labour markets (e.g. 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. Since the SCE re-interviews households for several months, we can see that a household's hiring rate expectations are very persistent.

We then show that a model in which households are rationally inattentive to the future hiring rate replicates both the heterogeneity and persistence of hiring rate expectations, and also preserves the possibility of multiple steady states through self-fulfilling expectations and precautionary saving.

Specifically, we assume that households cannot directly observe the hiring rate. Following Sims (2003) and others in the rational inattention literature, households can obtain signals about the hiring rate, but this information processing is costly. The cost is proportional to the informativeness of the signals, and so households choose to only process noisy signals before deciding on their consumption. The hiring rate is naturally bounded between 0 and 1, so the optimal signal structure chosen by households has a discrete number of 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. We also assume that households 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 very similar to what Adam and Marcet (2011) refer to as 'internal rationality'. To our knowledge, we are the first paper to examine this combination of rational inattention and incorrect prior beliefs. This is motivated by another observation from the SCE: 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 combination of information processing costs and imprecise prior beliefs is central to our results. In an unemployment trap steady state, most households receive signals that the hiring rate is low, and try to save as a result. The signal contains a significant amount of noise, so they do not know precisely how low the hiring rate is. This noise is particularly large when households have dispersed prior beliefs: in that case households have to use their signals to distinguish between a wide variety of possible states of the world. If instead households knew that the hiring rate had become concentrated around a low employment steady state, their prior beliefs would become much more precise. After collecting signals, they would have a very accurate expectation about the hiring rate, and so their behaviour would closely mimic optimal behaviour under full information. As 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. We would also not see the significant belief heterogeneity found in the data.

However, despite the importance of imprecise prior beliefs, multiplicity does not disappear over time as households observe a long history of signals and update their beliefs. In order to learn and update their beliefs about the distribution of the hiring rate, households must process new information, which comes with a cost. If a household begins in some period with a very dispersed prior belief about the distribution of the hiring rate, they will process some information, act on it, and update their beliefs for the next period. In the next period, the value of new information has been greatly reduced, because their prior belief already contains information gathered the previous period. The costs of new information, however, have not changed, and so the household chooses to process much less information in the second period. 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].

Hiring rate expectations are heterogeneous in our setup because households receive signals with idiosyncratic noise. Expectations are persistent because households update their beliefs about the hiring rate distribution each period, and believe that there is some persistence in that variable. If a household receives a signal that the hiring rate is low in one period, then they will update their beliefs about the distribution of that rate to put more weight on low hiring rates. For any given signal in a future period, they will therefore form a posterior belief which is biased downwards.

Section 2 places this work in the context of the literature. Section 3 explores the survey data

[1]In Adam and Marcet (2011), beliefs collapse to 'near-rationality', the point at which 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.

on hiring rate expectations. In section 4 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 5. We show that the combination of rational inattention and imprecise prior beliefs about the hiring rate generate two possible steady states: an intended steady state with a high hiring rate and moderate inflation, and an unemployment trap with a low hiring rate, and low (but positive) inflation. 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 exist even in a model with the heterogeneity in expectations which is a feature of the data. The co-ordination of beliefs which Morris and Shin (2000) argue is often necessary in models with multiple equilibria is not required in our model.

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

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

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. To our knowledge we are the first to combine these two information restrictions. We deviate somewhat 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 3 we document that the same is true for hiring rate expectations. The theoretical implications of heterogeneity in inflation expectations have been studied by Andrade et al (2019), Wiederholt (2017), among others. In contrast to this literature, we study heterogeneous labour market expectations.

3 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 variable of interest in Ravn and Sterk (2018). The figure below shows the histogram of responses, normalised by the average response for the month of the interview.

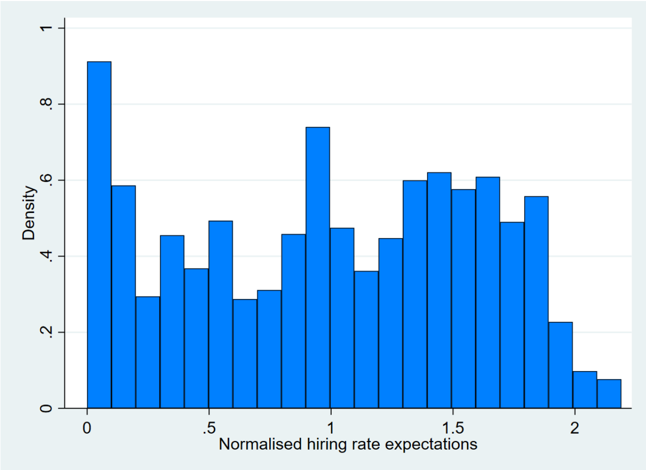

Figure 1: Histogram of hiring rate expectations from the Survey of Consumer Expectations, normalised by the average for the month of the interview.

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. Household beliefs about the economy could be co-ordinated, but that would imply different hiring rates for high and low education households, for example. To explore this, we run the following regression:

$$\mathbf{E}_{it}\eta_{it} = \beta_0 + \beta_1 X_{it} + \beta_2 W_t + \varepsilon_{it} \tag{1}$$

The dependent variable is household i's expectation of their own hiring rate. X_{it} collects every personal characteristic available in the SCE[3]. W_t are month fixed effects, to pick up any macroeconomic influences on the responses. The R^2 for this regression is 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.

The belief co-ordination required for existing theories of self-fulfilling expectation driven fluctuations is therefore implausible. Labour market expectations are heterogeneous. They are also persistent at the household level. Table 1 shows the results of regressing ε_{it}, the part of hiring rate expectations which cannot be explained by observed household characteristics or macroeconomic

[3]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.

conditions, on the lag of itself[4] $\varepsilon_{i,t-1}$. The coefficient is large and positive: a household which was optimistic one month is likely to remain optimistic in the next month.

Table 1: Regression of unexplained part of hiring rate expectations on the lag of itself

	(1)
	Hiring Rate Residual
L.Hiring Rate Residual	0.718***
	(0.00390)
Constant	-0.0279
	(0.118)
Observations	31830

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

One model consistent with these results is that households receive signals about the hiring rate which contain idiosyncratic noise, and they do not observe past values of the hiring rate, so they must rely on the history of signals to form their expectations. The rational inattention model we study fits into this class of models, and so will feature heterogeneity and persistence of expectations.

Our final observation from this data is that the variance of expectations is significantly 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[5] of this series at 4.7. In contrast, the standard deviation of the unexplained part of hiring rate expectations ϵ_{it} is 30.5. 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)), the standard deviation of expectations would be *lower* than the standard deviation of the data[6]. 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. This motivates our assumption that households hold 'ignorance priors', discussed further in section 4.7.

[4]This is possible because households remain in the SCE for several months.

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

[6]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.

4 Static Model

This section begins with a simple two period version of the consumer problem in Ravn and Sterk (2018). The firm side is also kept very simple, to illustrate the key forces driving the unemployment trap. In section 4.8 we construct a simple dynamic model by repeating the two period 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 as households learn over time.

4.1 Households

There is a unit mass of households. All households are employed in period 1, and they receive a wage of 1. With probability ω a household loses their job at the end of period 1. There is then a round of hiring by firms, so newly separated workers find employment for period 2 with probability η. If they are employed in period 2, 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 period 1 at interest rate R. The budget constraints for household i in period 1, period 2 if employed, and period 2 if unemployed (respectively) are:

$$c_{1i} + \frac{b_{1i}}{R} = 1 \tag{2}$$

$$c_{2i}^{e} = w + b_{1i} \tag{3}$$

$$c_{2i}^{u} = \theta + b_{1i} \tag{4}$$

Combining these we have:

$$c_{2i}^{e} = w + R(1 - c_{1i}) \tag{5}$$

$$c_{2i}^{u} = \theta + R(1 - c_{1i}) \tag{6}$$

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

The household problem is therefore to choose c_{1i} to maximise:

$$V_{1i} = -\frac{1}{2}(c_{1i} - \bar{c})^2 - \frac{1}{2}\beta \mathbf{E}_{1i}\omega(1-\eta)(c_{2i}^{u} - \bar{c})^2 - \frac{1}{2}\beta \mathbf{E}_{1i}(1 - \omega(1-\eta))(c_{2i}^{e} - \bar{c})^2 \tag{7}$$

If the households observe the hiring rate η all households make the same choice of c_1, to satisfy the FOC:

$$c_1 = \beta R\omega(1-\eta)\big(\theta + R(1-c_1)\big) + \beta R\big(1-\omega(1-\eta)\big)\big(w + R(1-c_1)\big) \tag{8}$$

In this case, period 1 consumption is an increasing linear function of the period 2 hiring rate, due to a simple precautionary savings motive. Solving out for c_1 we have:

$$c_1 = \frac{\beta R}{1+\beta R^2}\big(R + w - \omega(1-\eta)(w-\theta)\big) \tag{9}$$

4.2 Rational Inattention Problem

We now relax the assumption that households can precisely observe the 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 13 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 7. The agent views the hiring rate as exogenous and random[7]. 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.

To simplify this problem, we can note that the value function in equation 7 is strictly concave in c_{1i}, so there is a unique function mapping the optimal consumption choice c_{1i}^* to the expected hiring rate $\mathbf{E}\eta$. Furthermore, the optimal 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 hiring rate and optimal consumption. An agent 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

[7] 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 4.7

mapping between signal realizations and the optimal consumption choice. We can therefore leave the signal choice in the background of the problem, and instead study the optimal decision rule linking consumption to the hiring rate, subject to the information costs of implementing such a rule.

Specifically, we express the household's decision rule as the joint probability density function over c_{1i} and η. That is, given a particular hiring rate η, the household chooses how often they will choose each different possible level of 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_{1i} for each level of η, given the information costs of reducing those mistakes. This follows the form of the rational inattention problem faced by firms in Matejka (2017).

The household problem is therefore:

$$f = \arg\max_{\hat{f}} \mathbf{E}[V(\eta, c_1)] - \psi\kappa(\hat{f}, g) = \arg\max_{\hat{f}} \int_{\eta}\int_{c_1} V(\eta, c_1)\hat{f}(y, x)d\eta dc_1 - \psi\kappa(\hat{f}, g) \tag{10}$$

subject to

$$\int_{c_1} \hat{f}(\eta, c_1)dc_1 = g(\eta) \qquad \forall\eta \tag{11}$$

$$\hat{f}(\eta, c_1) \geq 0 \qquad \forall\eta, c_1 \tag{12}$$

$$\kappa(\hat{f}, g) = H[g(\eta)] - \mathbf{E}_{c_1} H[\hat{f}(\eta|c_1)] \tag{13}$$

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

$$H[g(\eta)] = -\int g(\eta)\log g(\eta)d\eta \tag{14}$$

The first constraint (equation 11) ensures that the marginal distribution of η obtained from the optimal joint pdf is consistent with $g(\eta)$, the household's prior belief about the distribution of the hiring rate [8].

The second constraint (equation 12) is that the solution must be positive everywhere, as required for it to be a joint pdf.

The final constraint (equation 13) is the information processing constraint. Entropy $H[.]$ is a

[8] 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 4.7).

measure of the dispersion of a distribution. The first term of constraint 13 is the entropy of the prior. The prior reflects the information held by the agent about the distribution of the hiring rate before receiving any signals. We will assume for now that this prior is uniformly distributed (this is relaxed in section 4.8), so the prior is rather dispersed and entropy is high. The second term is the expected entropy of $f(\eta|c_1)$, the updated distribution over η believed by the household after taking in the available signals[9]. A precise 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.

Given a particular level of κ, the agent 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 processing capacity. This means that the agent 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 $f(\eta|c_1)$ must remain dispersed. When η is in that small range, the household making that decision will be very accurate in choosing optimal c_1, 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 estimating η, 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.

4.3 Rational Inattention Solution

This problem follows the setup of the firm problem in Matejka (2017), in which firms optimally choose from a small number of discrete prices even though there is a continuum of prices available to them.

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 $\psi = 0.002$ is plotted in figure

[9]The 'information content' of the signals is incorporated into the choice of c_1, so the conditional distribution of η given the choice of c_1 tells us what the agent believes about η.

2. This 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 $\kappa \approx 0.5$. As in Matejka (2017), 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. The logic behind this is discussed in section 4.2 above, and in detail in Jung et al (2015) and Matejka (2017). As η increases (and so the optimal choice of c_1 under perfect information increases), the probability an agent chooses the higher level of consumption in their 'menu' increases.

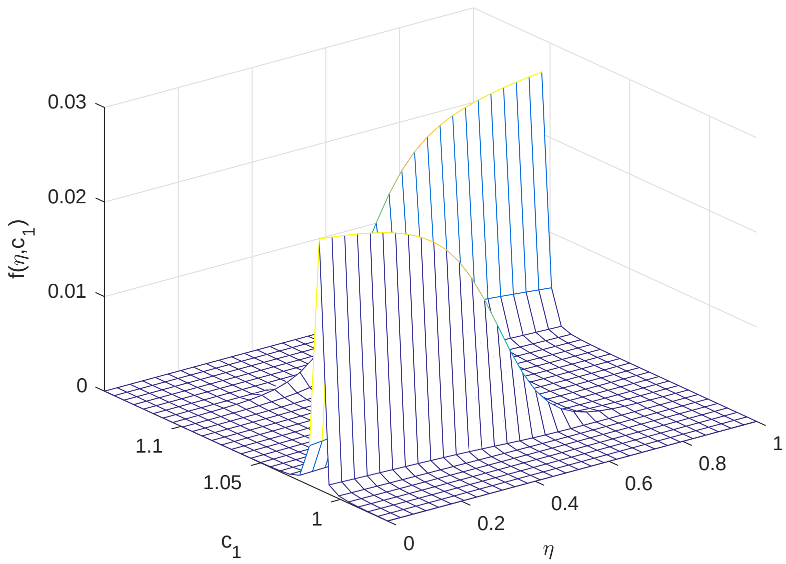

Figure 2: Optimal decision rule for $\psi = 0.002$

4.4 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 consumption levels in the optimal 'menu' which arises from the RI problem with uniform priors (equations 10 - 13). The proportions on each level of consumption are determined by the probabilities in the optimal joint pdf obtained as the decision rule from that agent problem.

Therefore for each level of η we obtain aggregate c_1 using:

$$\bar{c_1}(\eta) = \int_{-\infty}^{\infty} c_1 f(c_1|\eta) dc_1 \tag{15}$$

With no information processing[10], households choose consumption to maximise $\mathbf{E}V$ given their prior belief (where $\mathbf{E}\eta = 0.5$), 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_1 for different η. For low values of κ, agents optimally restrict themselves to two values of consumption. Importantly, the aggregate consumption function has a wave-like shape around its perfect information equivalent.

As shown in Matejka (2017), as κ rises further, more choices of c_1 are introduced into the optimal menu. As this occurs the aggregate response of 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 B.1.

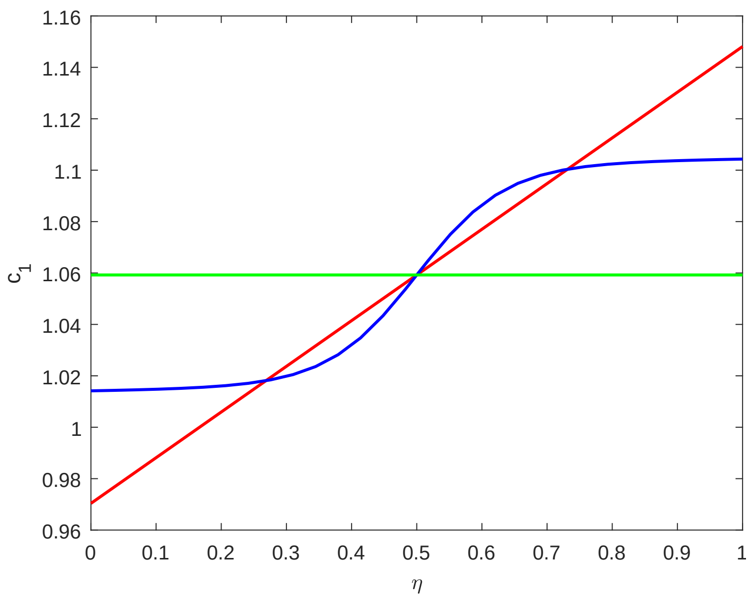

Figure 3: Aggregate consumption function with $\kappa = 0$ (green), $\kappa = 0.5$ (blue) and in the unconstrained case (red)

Consider the case where information processing is constrained but non-zero (the blue curve). This meets the full information consumption function (in red) at $\bar{c}_1 = 1.06, \eta = 0.5$. However, at this hiring rate in the full information model, every household consumes $c_1 = 1.06$. In contrast, in the rational inattention model, half of the households get a signal that the hiring rate is 'high' and choose $c_1 = 1.104$ accordingly. The other half receive a signal that η is 'low', and so consume

[10] Any $\psi > 0.1$ leads to ($\kappa = 0$) being optimal.

$c_1 = 1.014$.

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 2 above, it can be seen that this is the case for extreme high and low values of η, and $\bar{c_1}(\eta)$ is flat in these regions accordingly. In contrast, as η moves from 0.4 to 0.6, large numbers of agents switch from choosing the low level of consumption to the high level, and this corresponds to the steep section of the corresponding aggregate response function in figure 3.

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 2 fell linearly as η increased the aggregate consumption function would be linear. In fact, the distribution of $c_1|\eta$ for the values of c_1 in the optimal menu is logistic in shape, which is what gives rise to the wave-like shape of the aggregate response curve[11]. 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.

4.5 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[12]. For the purposes of this simple two period example, it will be sufficient to say that firms hire more workers when aggregate demand is high, so:

$$\eta = H(\bar{c}_1), \text{ with } H'(\bar{c}_1) > 0 \tag{16}$$

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 A microfounds such a process for the hiring rate using a model with working capital. In section 5 the firm side of the model is standard, as in Ravn and Sterk (2018).

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

[12] 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.

4.6 Equilibrium

It is important here that the hiring rate η is not in fact uniformly distributed like the prior beliefs of the agents. In Matejka (2017) and many other models in the rational inattention literature, agents' prior beliefs about the distribution of the unknown variable are correct. However, this may not be a plausible assumption for 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. The survey data in section 3 suggests that for the hiring rate, households do not know the true equilibrium distribution. We therefore explore the situation where households have priors that depart from reality. This might be particularly pertinent, for example, after structural breaks or times of turmoil. This assumption is very similar to the 'internal rationality' in Adam and Marcet (2011).

In particular, we assume households start in some initial period from 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[13]. This is discussed in section 4.7. We begin by studying this initial period. In section 4.8 we show that the multiplicity does not disappear when agents update their priors away from this uniform starting point. The separation of the prior belief and the true equilibrium distribution of η means that the aggregate consumption function $\bar{c_1}(\eta)$ remains as in figure 3, and the hiring rate is then determined endogenously by the interaction of this aggregate response and the firm hiring function 16.

The graph below shows this equilibrium interaction. The blue and red aggregate consumption functions are as in figure 3. The equilibrium condition 16 is added in black. Under full information there is one equilibrium, but under RI there is also an unemployment trap.

[13]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.

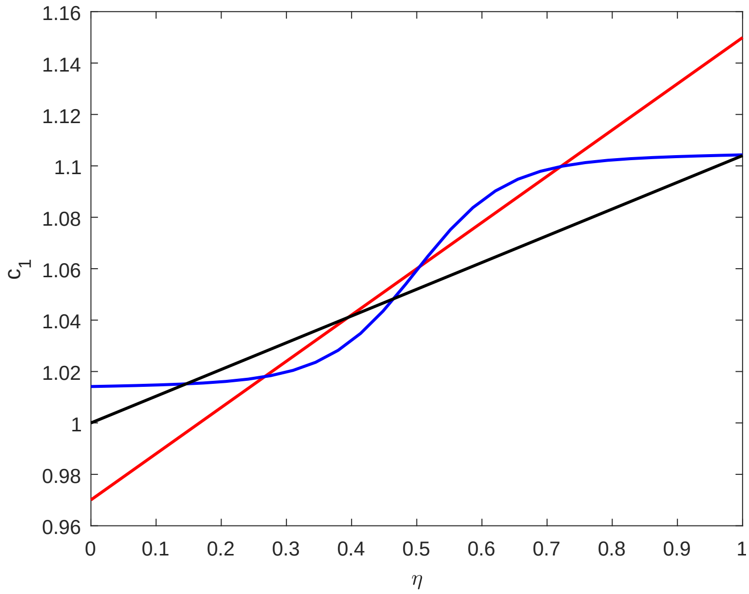

Figure 4: Aggregate consumption response to changes in the hiring rate for full information (red) and rational inattention with $\psi = 0.002$ (blue), with the firm condition 16 in black

Consider first the rational inattention equilibrium close to $\eta = 0.5$. As explained in section 4.4, at this hiring rate 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 $c_1 = 1.104$ to $c_1 = 1.014$. This means that aggregate consumption falls a great deal, so much so that there is another equilibrium at $(\eta = 0.144, c_1 = 1.015)$ in which almost all agents 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 c_1 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. A similar process gives rise to the equilibrium close to $\eta = 1$.

In this simple model, the middle equilibrium will feature heterogeneous hiring rate expectations among households, but the unemployment trap features near-universal agreement. Almost all households agree that the hiring rate is low, and they agree on how low that is because they have all adopted the same information strategy. This does not mean, however, that we cannot generate multiple equilibria which all feature expectations heterogeneity: if the information cost ψ was lower, households would process more informative signals, and would have more consumption levels in their optimal menu. The smoothed step shape of the consumption function would remain, but there would be more 'steps'. We would therefore see multiple equilibria which feature significant belief heterogeneity.

4.7 Conditions 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 must not know the true equilibrium distribution of the hiring rate.

To see why, consider the equilibrium in figure 4. The true equilibrium distribution of η contains just three 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, as they do not need to worry about the possibility of an extremely low η. This, in turn, would mean that the equilibrium values of η would be closer to 0.5. 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 'ignorance prior' assumption is very similar to the removal of 'external rationality' in Adam and Marcet (2011). Like them, we believe that it is implausible to suppose that agents know the true stochastic processes and mechanisms underlying the determination of the endogenous variables they face, especially when those mechanisms involve the choice behaviour of many other agents[14]. The survey data on hiring rate expectations in section 3 supports this view.

In the HANK model in section 5, we only consider the stationary equilibrium. It is, however,

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

based on the model in Ravn and Sterk (2018) which contains a variety of shocks. These affect policy variables, household decisions, and firm profit functions. The decision of how many vacancies a firm will post therefore depends on a variety of shocks, and on how a variety of agents in the economy respond to those shocks. The same is true of the unemployment rate. As an example, consider a shock to technology. The shock directly affects the firm profit function through marginal costs. Firms therefore change their vacancy posting decisions. There are also important general equilibrium effects. Households change their consumption demand, and firms also change prices. This will lead to a response of the interest rate through the Taylor Rule. Furthermore, the aggregate vacancy decisions of the population of firms will affect labour market tightness, which affects the returns to posting a vacancy. The mechanisms linking a technology shock to the hiring rate are therefore complicated, and rely on multiple endogenous reactions from different economic agents.

The households in most standard rational inattention papers (e.g. Mackowiak and Widerholt, 2015) do not have ignorance priors, because they have a complete understanding of how the model operates. They process costly information about the technology shock[15], and so decide consumption, because they know how the technology shock propagates through the economy to the hiring rate. Households understand all of the links outlined in the previous paragraph. They have common knowledge of how all other households make their information and consumption decisions, as this will be part of how the shock affects the hiring rate. This kind of model therefore assumes that households have a large amount of information about the economy and about how other households make decisions, but then be places limits on the information that those very well-informed households can obtain about the realisations of the shocks.

A more realistic set of assumptions is that households do not perfectly understand the links from shocks to aggregate household choices to the hiring rate. In this case we have an ignorance prior: households have a subjective belief about the endogenous state variable that forms their prior, and they process information to improve these beliefs, but that prior is not tied to the true equilibrium distribution of the state variable.

The other way in which households could have accurate prior beliefs about the hiring rate would be if they had observed the hiring rate over many periods, and had learned its equilibrium distribution over time. This is similar to the assumption in Adam and Marcet (2011) that agents are

[15]In addition, households know the true distribution of the shocks.

'near-rational', which they interpret to mean that their 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. However, in section 4.8 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. This is because the initial reduction in dispersion of beliefs after one period of information processing leads agents to process less information in the following periods, and so prior beliefs do not approach the true distribution of the endogenous variable. This suggests that the assumption in Adam and Marcet that agents get close to the true equilibrium distribution of endogenous variables is implausible in models where processing information is costly.[16] Therefore the ignorance prior is a plausible assumption for household beliefs about the hiring rate.

The other important condition for multiplicity is that the firm hiring function $H(\bar{c}_1)$ 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 $\frac{dc_1}{d\eta} > 0$. The rational inattention aggregate consumption function cannot therefore be downward sloping. This means that if $H'(\bar{c}_1) < 0$, there will be just one equilibrium.

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

[16] 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.

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 5 we find an unemployment trap with low, but positive, labour market activity, whereas the unemployment trap in Ravn and Sterk has zero employment.

4.8 Dynamic Solution

When Adam and Marcet (2011) discuss models with internal rationality they make a similar assumption to us, that agents do not know the true distribution of variables in their environment. However, they consider situations where agent prior beliefs about those variables are close to the true distributions. This is because they allow their agents to observe data and learn, so they would discard any beliefs which clearly do not generate the data. The model in section 4.1 is static, so there is no room for agents to learn in this way. In this section, we repeat the two periods from that simple model to give households the opportunity to learn about the distribution of the hiring rate. For this section, a 'period' refers to an iteration of the original two-period model. 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. We show that in this simple model, beliefs do not converge to the true distribution of η, even in the long run. The multiplicity therefore holds even in steady state.

We will consider a very simple rule for updating beliefs. If the household reaches the posterior belief $f(\eta_t|c_t)$ in period t, we will suppose that their prior belief in period $(t+1)$ is given by:

$$g(\eta_{t+1}|c_t) = \rho f(\eta_t|c_t) + (1-\rho), \quad \rho \in (0,1) \tag{17}$$

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[17].

We will also make one further simplification, which is to assume that when households make

[17]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.

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 full rational model where future information values are taken into account.

With this assumption, the first period optimal $f(\eta_1, c_1)$ (from the uniform prior) is as in figure 2. 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 4.2. There are two possible posterior beliefs that the households could finish period 1 with. These are plotted below.

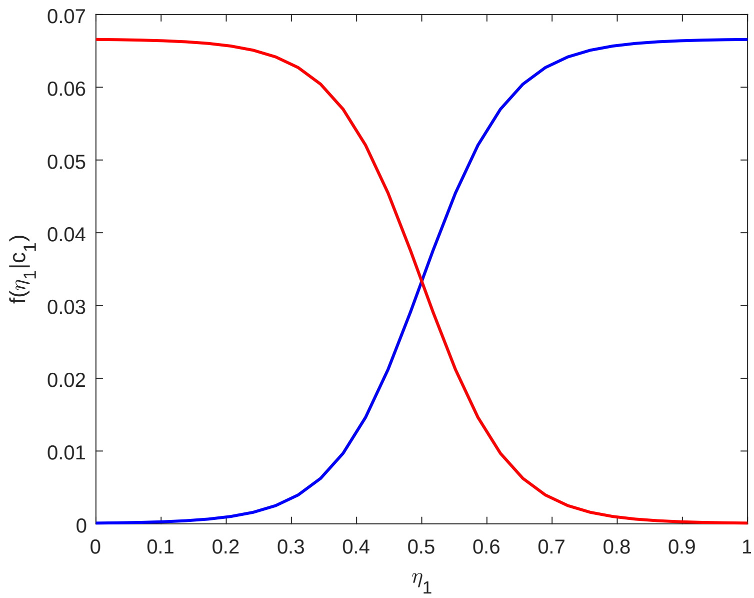

Figure 5: Posterior beliefs from a uniform prior and $\psi = 0.002$ for those who chose low c (blue) and high c (red)

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 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 c_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 c_2 than an agent who chose a higher c 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 $\rho < 1$, period 2 priors are more dispersed than period 1 posteriors[18], though they remain 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 agents 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 difference between our model and that of Adam and Marcet is that the only limit on learning in their model is the availability of data. Here, 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.

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 $\bar{c}_1$ 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

[18]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.

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. 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 these candidate steady states. The graph below plots these candidate steady states when $\rho = 0.9$: the $\bar{c}$ that comes out of this iteration process for each possible $\bar{\eta}$. Again, the equilibrium condition from the firm side of the model 16 is in black, and the model has steady states where the blue set of candidate steady states meets this equilibrium condition.

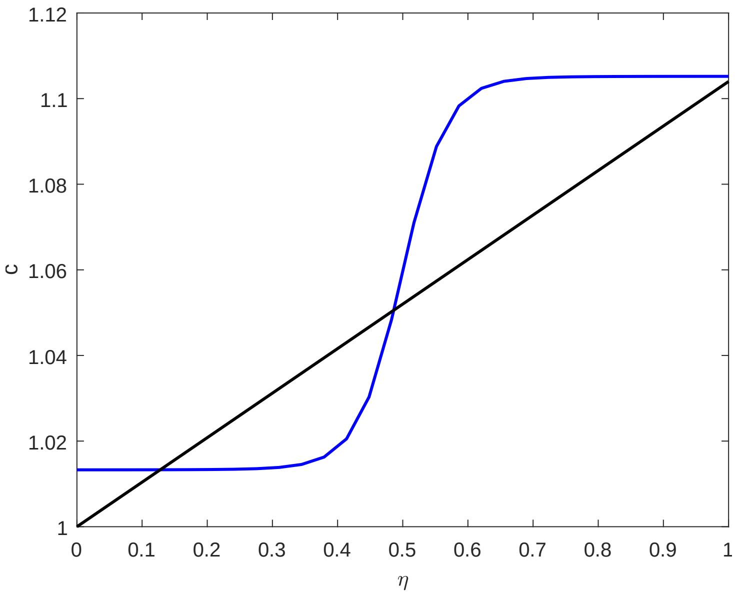

Figure 6: Candidate steady state $\bar{c}$ for each η.

If beliefs did degenerate to the truth here, this curve would coincide with the linear full information response curve in figure 3. 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 an intended steady state with $\eta = 0.48$, and an unemployment trap steady state with $\eta = 0.13$. Note that the blue curve here is not a consumption function with the same meaning as figure 3. It plots out potential steady states, in which the distribution of beliefs, and so of consumption, is stable if η remains stable. Each of those potential steady states has an associated aggregate consumption function, determined by the mix of prior beliefs present at the steady state. For example, for the unemployment trap steady state at $\bar{\eta} = 0.125$, the aggregate consumption function is plotted below.

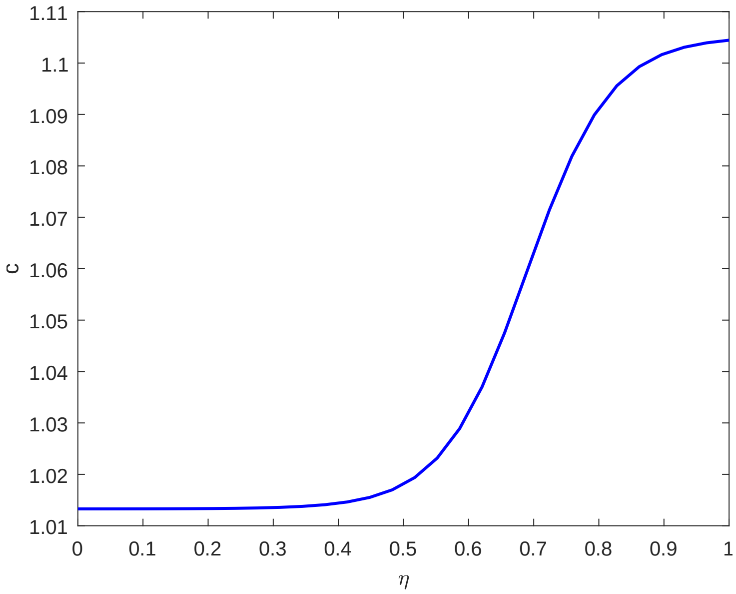

Figure 7: Aggregate response of c to η if belief composition is as in the unemployment trap steady state.

When η has been at 0.125 for many periods, most households are choosing a low c each period, and so most have priors which are biased towards low c. For significant numbers of them to choose a higher c, as needed for aggregate c to rise, η must therefore be much higher than is necessary when every agent has a uniform prior as in period 1. The equivalent aggregate consumption function from the intended steady state is displayed in appendix B.2.

Interestingly, even when η is held close to 0 for many periods, no household chooses c below 1.014, the low level of c in the period 1 optimal menu. If a household chooses $c_t = 1.014$, their priors in the next period will bias them towards making that same choice again. We have established that the household will not process much information in period $(t+1)$ because their priors already contain lots of information, so the marginal value of more information is low. We have also argued, in giving intuition for the discreteness result in section 4.3, that households try to avoid large mistakes when they choose the structure of their information. The small amount of information the household processes in period $(t+1)$ will therefore be structured to help the household avoid large mistakes first, with learning about finer detail of η a second-order concern. Receiving a signal if η is in fact very high helps the agent avoid the large mistake of choosing a low c when η is very high. This is much more valuable for expected utility than distinguishing between 'very low' and 'somewhat low' η, so the household never processes the information required to adjust down from their low prior, and so the lowest value in the household's optimal menu of choices remains $c = 1.014$. This argument applies in reverse at the high end of the η range: no household ever chooses $c > 1.104$, because to gather the information needed to do so would not be the best use of information processing capacity.

5 HANK model

Here we study the HANK model in Ravn and Sterk (2018) 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 only small adjustments to this 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.

5.1 Firms

Firms set prices and choose vacancies to maximise profits[19]. There are quadratic adjustment costs as in Rotemberg (1982), firms are monopolistic and households have CES preferences. The labour market matching function is:

$$M(u_s, v_s) = u_s^{\alpha} v_s^{1-\alpha} \tag{18}$$

We assume that wages are set exogenously, and are independent of the hiring rate. They are set such that they at least compensate the worker for the disutility of labour. Using these assumptions and functional forms, the Phillips curve becomes:

$$1 - \gamma + \gamma mc_s = \phi(\Pi_s - 1)\Pi_s - \phi \mathbf{E}_s \Lambda_{s,s+1} \left[\frac{y_{s+1}}{y_s} (\Pi_{s+1} - 1)\Pi_{s+1} \right] \tag{19}$$

where

$$mc_s = \frac{w_s}{e^{A_s}} + \frac{k}{q_s} - (1 - \omega)\mathbf{E}_s \Lambda_{s,s+1} \frac{k}{q_{s+1}} \tag{20}$$

In steady state this becomes:

$$\phi(1 - \beta)(\Pi - 1)\Pi = 1 - \gamma + \gamma(w + k\eta^{\frac{\alpha}{1-\alpha}}(1 - \beta(1 - \omega))) \tag{21}$$

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 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.

Equation 21 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 will only study the region of the PC where vacancies (and so hiring) are strictly positive, so λ_f is always zero in this version of the model.

[19] This part of the model is identical to that in RS. Their paper sets out the firm problem in detail.

5.2 Households

The employed household's problem is as in RS:

$$\max V_s^e = \Big(\frac{c_{es}^{1-\mu}-1}{1-\mu} - \zeta\Big) + \beta \mathbf{E}_s \omega(1-\eta_{s+1})V_{s+1}^u + \beta \mathbf{E}_s(1-\omega(1-\eta_{s+1}))V_{s+1}^e \tag{22}$$

subject to

$$P_s c_{es} + \frac{b_{s+1}}{R_s} = W + b_s \tag{23}$$

$$b_{s+1} \geq 0 \tag{24}$$

We will study the steady states of this model. In RS, the only unknowns are shocks. Since we set these all to zero, in the full information case there is no uncertainty for households except over their future employment status. We can therefore drop the expectation operators in the full information case.

RS note that the no-borrowing constraint is never binding for these households, so we can drop constraint 24. We express nominal wage W as the (fixed) real wage multiplied by the price level wP_s.

The first order condition under perfect information is as in RS:

$$c_{es}^{-\mu} = \beta \frac{R_s}{\Pi_{s+1}} \Big(\omega(1-\eta_{s+1})c_{u,s+1}^{-\mu} + (1-\omega(1-\eta_{s+1}))c_{e,s+1}^{-\mu}\Big) \tag{25}$$

Unemployed households are always at their borrowing constraint, and so their problem never matters for equilibrium determination[20].

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[21], and they substitute for an interest rate Taylor rule:

$$R_s = \max\{\bar{R}\bar{\Pi}^{-\delta_\pi}\Pi^{\delta_\pi}\eta^{\frac{\delta_\eta}{1-\alpha}}, 1\} \tag{26}$$

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

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

This gives:

$$1 = \frac{\beta \max\{R^*\Pi^{\delta_\pi}, 1\}}{\Pi}\left[\omega(1-\eta)\left(\frac{\vartheta}{w}\right)^{-\mu} + 1 - \omega(1-\eta)\right] \tag{27}$$

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

5.3 Perfect Information Steady State

This is the same as in RS, except that we are only considering the case in which $R > 1$, and we only study the region where $\eta > 0$. The Phillips Curve and Euler Equation for employed households are shown below in blue and red respectively[22]. The y axis is steady state inflation and the x axis is the steady state hiring rate.

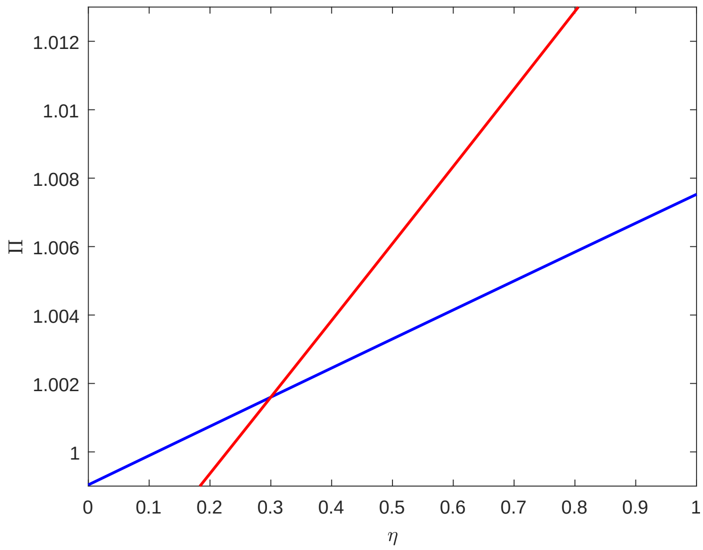

Figure 8: Steady State relations under perfect information. The Phillips Curve is in blue, the steady state Euler equation is in red.

There is only one steady state with positive hiring[23]. This figure is identical to figure 3 panel 1 in RS, with the exception that we have not plotted the steady state relationships at the zero lower bound as we do not study that region in this paper. Specifically, this is the outcome of the RS model when labour income risk is countercyclical, which is the case here because we assume that wages are independent of the hiring rate.

The 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

[22] Appendix C has the parameter values under which this is drawn. The calibration is taken from appendix A3 of Ravn and Sterk (2018).

[23] In RS there is an unemployment trap because the Phillips Curve is kinked, turning vertical at $\eta = 0$ due to the constraint that vacancies cannot be negative. This gives rise to a steady state with zero vacancy posting and so zero hiring and employment. With RI we find multiple steady states with $\eta > 0$.

steady state hiring rate η decreases the desire for precautionary saving. 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.

5.4 Rational Inattention

As in section 4.2, we now amend the household problem so that processing information about the future hiring rate has marginal cost ψ. Again, we start with uniform prior beliefs $g(\eta) \sim U(0,1)$ for all households. In section 5.5 we allow households to update their prior beliefs using information processed in the previous period, as in section 4.8.

Substituting the budget constraint 23 into the value function, the employed household problem with RI is as follows:

$$f = \arg\max_{\hat{f}} \mathbf{E}[V_s^e(\eta_{s+1}, b_{s+1})] - \psi\kappa$$
$$= \arg\max_{\hat{f}} \int_{\eta_{s+1}} \int_{b_{s+1}} V_s^e(\eta_{s+1}, b_{s+1})\hat{f}(\eta_{s+1}, b_{s+1}) d\eta_{s+1} db_{s+1} - \psi\kappa \quad (28)$$

subject to

$$\int_x \hat{f}(\eta_{s+1}, b_{s+1})dx = g(\eta_{s+1}) \qquad \forall \eta_{s+1} \quad (29)$$

$$\hat{f}(\eta_{s+1}, b_{s+1}) \geq 0 \qquad \forall \eta_{s+1}, b_{s+1} \quad (30)$$

$$H[g(\eta_{s+1})] - \mathbf{E}_{b_{s+1}} H[\hat{f}(\eta_{s+1}|b_{s+1})] \leq \kappa \quad (31)$$

The function $H[.]$ is the entropy of the distribution over which it operates. It is defined in equation 14 above. These constraints are explained in section 4.2. 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 all relevant parameters in the problem. 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 26) 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.

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[24]. 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 household problem under these assumptions for $\Pi = 1.0055$ and $b_s = 0$ is plotted[25] in figure 9. It shows how the savings choices of employed households with uniform prior beliefs vary with η_{s+1}.

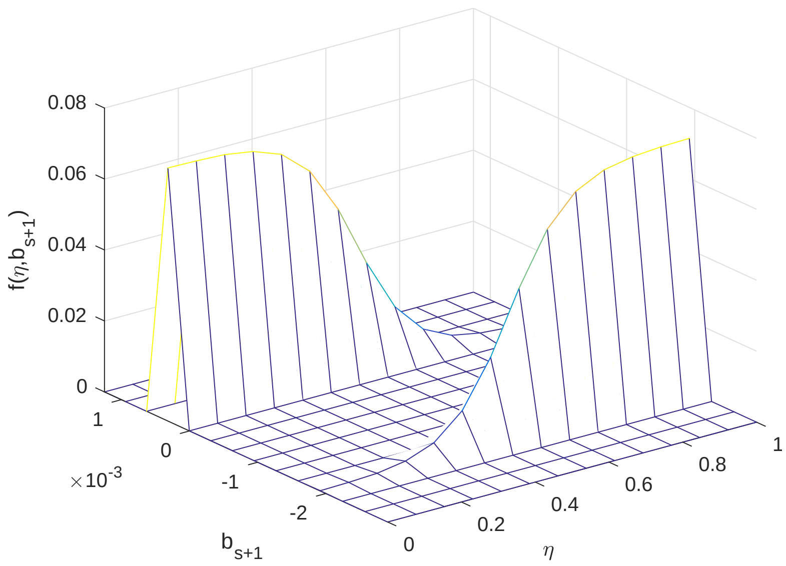

Figure 9: Decision rule for $\Pi = 1.0055$, $b_s = 0$, with a uniform prior belief.

As in section 4.3, households choose to limit themselves to a discrete number of savings choices. 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

[24]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 13

[25]All of the diagrams in this section are drawn for a monthly calibration. $\Pi = 1.0055$ is the inflation rate at which the full information Euler equation is in steady state with $\eta = 0.5$.

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.

5.5 Belief updating

As in section 4.8 for the simple model, we now allow households to update their beliefs over time using the simple rule:

$$g(\eta_{t+1}|c_t) = \rho f(\eta_t|c_t) + (1-\rho), \quad \rho \in (0,1) \tag{32}$$

A household with a uniform prior and zero wealth in period s (their decision rule is plotted as an example in figure 9) will therefore start period $(s+1)$ with one of two prior beliefs, plotted below. I 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).

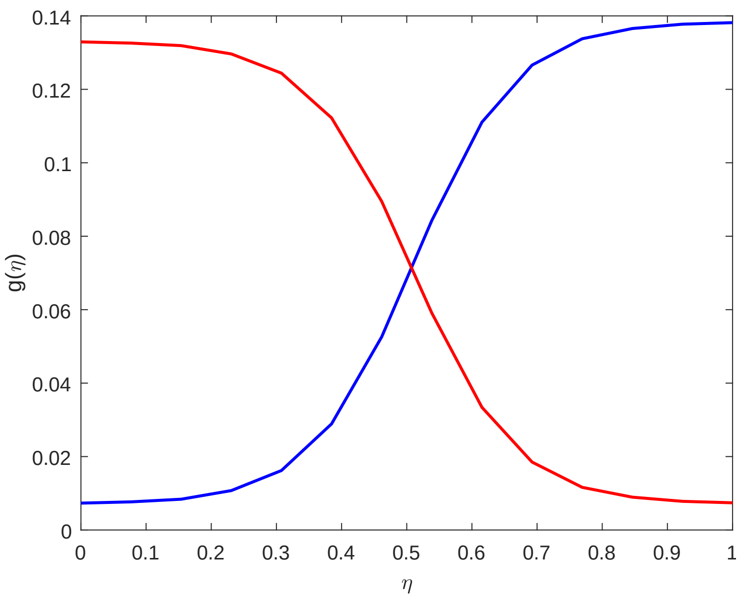

Figure 10: Possible prior beliefs after information processing from a uniform prior.

Starting from these two prior beliefs, the decision rules of a household with $b_s = 0$ are plotted below.

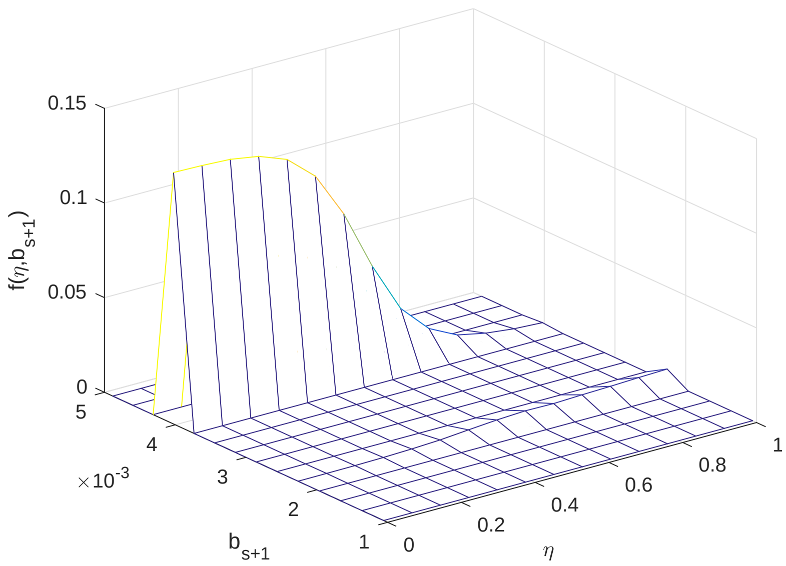

Figure 11: Decision rule for $\Pi = 1.0055$, $b_s = 0$, with a pessimistic prior belief.

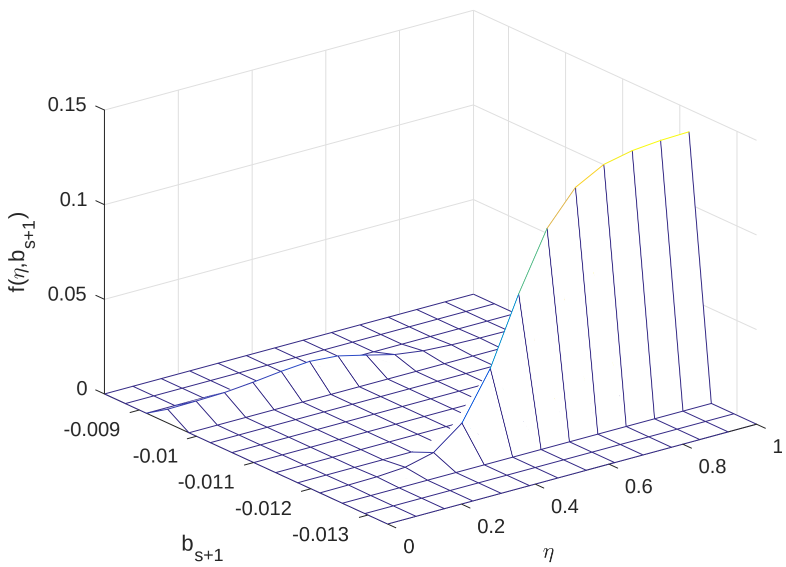

Figure 12: Decision rule for $\Pi = 1.0055$, $b_s = 0$, with an optimistic prior belief.

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 10: in steady state households are divided into optimists and pessimists.

5.6 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. Combining the decision rule in figure

11 and the pessimistic prior belief from figure 10 gives us that when $\eta = 0.5$, a pessimistic household with zero wealth will choose to save $b_{s+1} = 4\text{x}10^{-3}$ with probability 0.97, and $b_{s+1} = 2.3\text{x}10^{-3}$ with probability 0.03. If they choose $b_{s+1} = 2.3\text{x}10^{-3}$, it is because they have received a signal that η is high, and so in the next period they will have a more optimistic prior belief. This is possible in steady state as long as shifts in the wealth and beliefs of other households ensure that the distributions of wealth and beliefs are constant.

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 scale on the optimal savings choice (b_{s+1}) axis has been normalised so that full information optimal savings lie between 0 and 1.

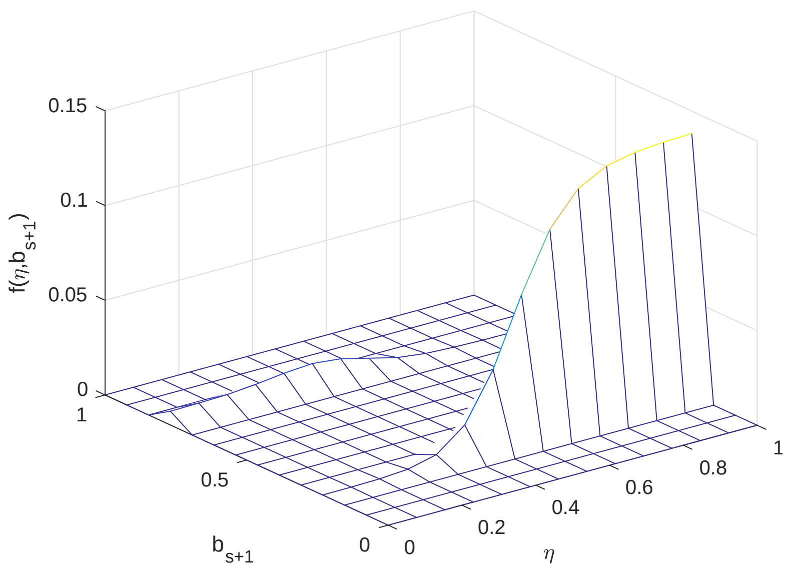

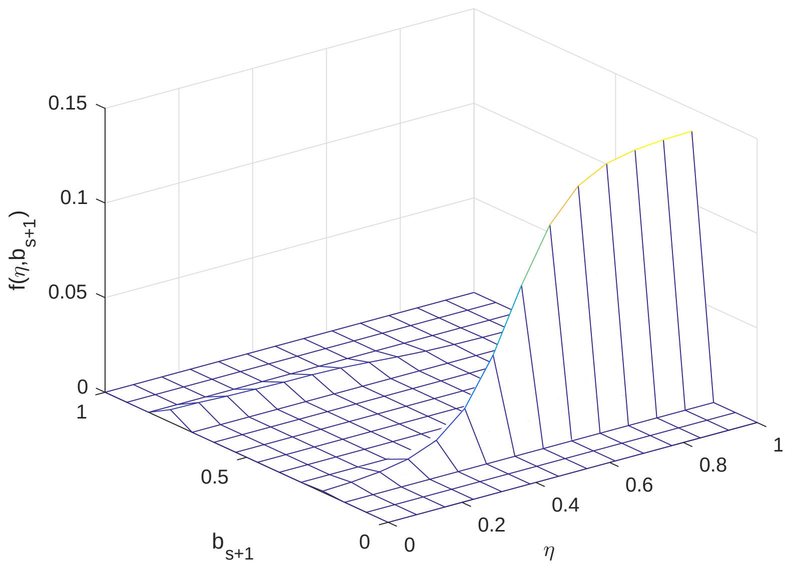

The first decision rule is for $b_s = 0$, and the second is for $b_s = 0.014$, which is the highest wealth seen in the steady states we compute. Without normalisation these two decision rules would have different scales on the b_{s+1} axis. The richer household chooses to save more than the poorer household 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 it makes the steady states significantly easier to find. 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 households with $b_s = 0$ - independently of the wealth distribution. After finding the steady state belief distribution, we then use this to find the steady state wealth distribution. This is significantly easier than trying to jointly determine the two distributions.

The approximation 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 X and Y.

The steady states of the model are plotted below. The blue and red lines are the Phillips Curve and full information Euler Equation from figure 8. 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.

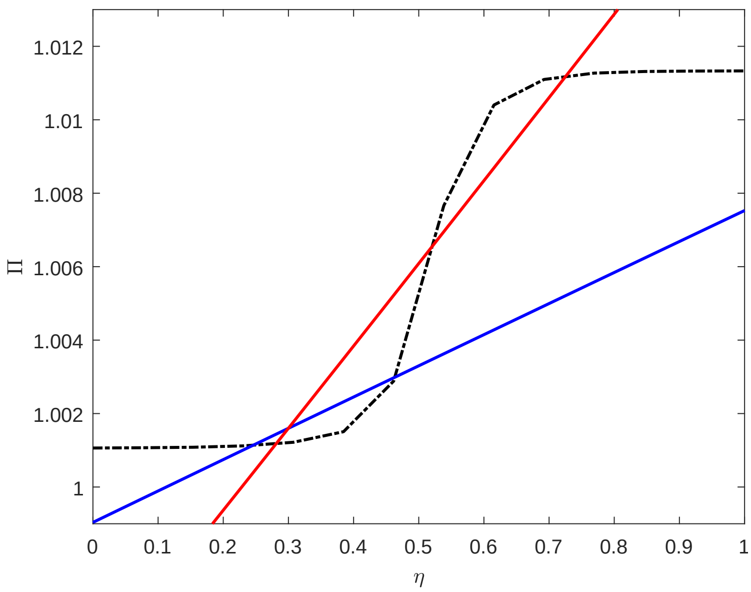

Figure 13: Steady state relations under perfect information and rational inattention ($\psi = 0.0025$)

Under perfect information there is a unique steady state, with $\eta = 0.3$ and $\Pi = 1.00165$. With 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 4.3. 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 9) therefore drives the multiplicity.

As in our simple example, the extreme steady states on the flat portions of the steady state Euler equation feature homogeneous prior beliefs across the population, and homogeneous choices. There is therefore very little belief or wealth inequality in those steady states. In the middle steady state, there is a significant variance in beliefs, and so in wealth.

As discussed in section 5.5, there are two groups of prior beliefs in steady state. This is a result of our choice of the information cost parameter ψ. The graph below plots the 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 10. 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.

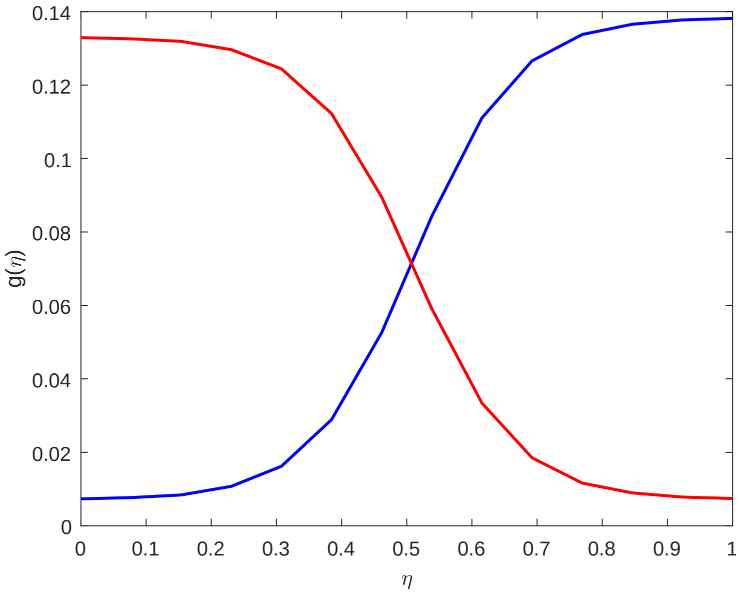

Figure 14: Steady state relations under perfect information and rational inattention ($\psi = 0.0025$)

There is churn underlying the steady states in figure 13. 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 for the unemployment trap and the intended steady state.

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

	$\mathbf{P}_{s+1}$	$\mathbf{O}_{s+1}$
$\mathbf{P}_s$	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 of households on pessimistic prior beliefs, while the high employment steady state has much greater dispersion of beliefs.

6 Conclusion

We have proposed a model of a self-fulfilling expectation driven unemployment trap which is consistent with key properties of survey expectations: labour market beliefs are heterogeneous, persistent, and display greater variance than that implied by accurate prior beliefs. The multiplicity of steady states is driven by households who face costs of processing information 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) adding rational inattention of this kind generates two steady states: a high employment steady state and an unemployment trap with low, but positive, inflation and hiring activity.

An interesting lesson from this analysis is that models with information processing costs can behave significantly differently when we drop the usual assumption that agents know the true marginal distributions of the variables of interest. The data from the Survey of Consumer Expectations suggests that this is the empirically relevant case for several variables; it would be interesting to see how this observation affects other applications of rational inattention.

Furthermore, we restrict ourselves to studying the steady states of the HANK model in section 5. It would be very interesting to study the dynamics of this model, to understand the policies or shocks that could cause the economy to transition between steady states in this environment of costly information processing.

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 Static Model: Firms

Here we derive the linear equilibrium relationship between aggregate consumption and the hiring rate in section 4.1.

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:

$$h = mv^{1-\alpha}u^{\alpha} \tag{33}$$

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[26]:

$$C(\eta) = km^{\frac{-1}{1-\alpha}}\omega\eta^{\frac{1}{1-\alpha}} \tag{34}$$

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:

$$D(\eta) = F - w - C(\eta) \tag{35}$$

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

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

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:

$$\bar{c}_1 - 1 = km^{\frac{-1}{1-\alpha}} \omega \eta^{\frac{1}{1-\alpha}} \tag{36}$$

In the graphs in section 4.1 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:

$$\bar{c}_1 = 1 + \frac{k\omega}{m}\eta \tag{37}$$

This is the linear relationship used in section 4.1. The calibration used there is approximately annual, with $\beta = 0.975$, $R = \frac{1}{\beta}$, $\omega = 0.4$, $m = 0.25$, $k = 0.05w$, $w = 1.3$, $\theta = 0.4$

B Aggregate consumption function graphs for the static model

B.1 Aggregate consumption function as cost of information falls

The graph below is the same as 3, 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 $\kappa \approx 1$.

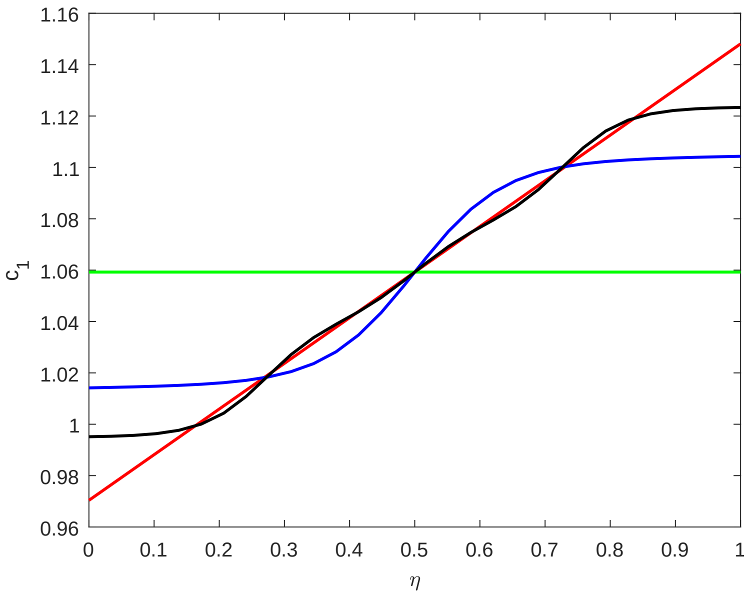

Figure 15: Aggregate consumption function with $\kappa = 0$ (green), $\kappa = 0.5$ (blue), $\kappa = 1$ (black) and in the unconstrained case (red)

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.2 Aggregate consumption function in updating case

The graph below is the same as figure 7, but with the addition of the aggregate consumption function derived under beliefs as in the intended steady state ($\eta = 0.48$) plotted in red.

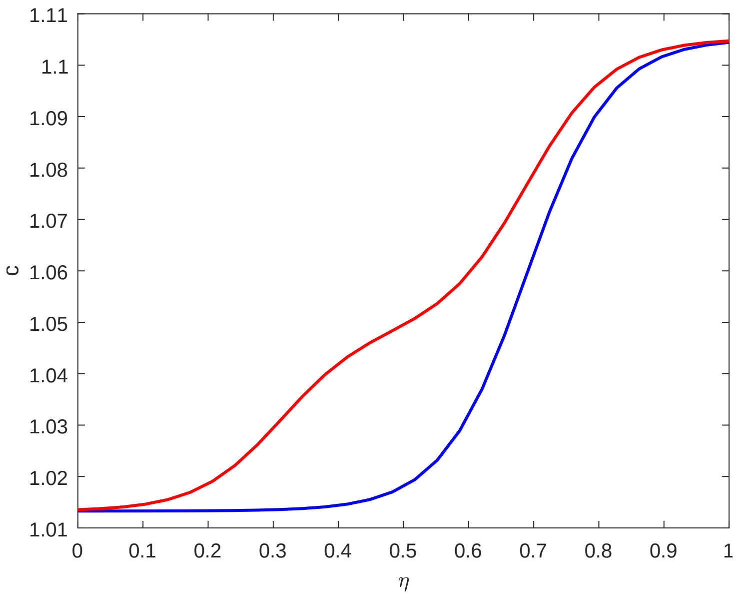

Figure 16: Aggregate response of c to η if belief composition is as in the unemployment trap steady state (blue) and the intended steady state (red).

The aggregate consumption function looks somewhat different at the intended steady state. The reason for the two steps is that, at this steady state, beliefs are not polarised one way or the other as in the unemployment trap. There is a mix of optimists (who have previously chosen high consumption and so are biased to believe that the hiring rate is high) and pessimists (whose beliefs are biased towards low η). As the hiring rate rises from zero, the optimists begin to believe that η is high, choose the high level of consumption, and so aggregate consumption increases. Once η has reached 0.4, almost all of the optimists are receiving signals that the hiring rate is high, and consuming accordingly. The pessimists, however, almost all get signals that the hiring rate is low at this point, and that continues to be the case as the hiring rate rises past 0.5. That is why the aggregate consumption function is flatter in this region. Beyond $\eta = 0.6$, as η rises the pessimists start to get signals that the hiring rate is in fact high, and so they increase their consumption in that range. The mixture of beliefs gives rise to the three-step shape.

C HANK model calibration

These parameter values are used to draw the figures in section 5. 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