

A Rational Inattention Unemployment Trap*

Martin Ellison[†] and Alistair Macaulay[‡]

August 11, 2021

Abstract

We show that introducing rational inattention into a model with uninsurable unemployment risk can generate multiple steady states, when the model with full information has a unique steady state. The model features persistent, heterogeneous labour market expectations, consistent with survey evidence. In a heterogeneous agent New Keynesian model, rational inattention to the future hiring rate generates three steady states: an unemployment trap with mild deflation and a low (but positive) job hiring rate, a middle steady state with moderate employment and inflation, and an ‘employment trap’ with high employment and inflation. Large mutations in the distribution of household beliefs can shift the economy between steady states.

*We are grateful to the editor and three anonymous referees for their extensive feedback. We are particularly indebted to the referee who pointed us to the proof in Appendix D.1. We thank Paul Beaudry, Gianluca Femminis, Alexandre N. Kohlhas, Albert Marcet, Filip Matějka, Edoardo Palombo, Morten Ravn, Ricardo Reis, Thomas J. Sargent, Alireza Sepahsafari, Christopher Sims, Mirko Wiederholt, and participants at various seminars and conferences for helpful comments.

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

[‡]Department of Economics & St. Anne’s College, University of Oxford. Email: alistair.macaulay@economics.ox.ac.uk

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). An important example concerns 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 (Carroll and Dunn, 1997). The resulting fall in aggregate demand causes employment to fall, confirming the pessimistic beliefs. The sharp rise in the savings of wealth-poor households, who are least able to self-insure against unemployment risk, suggests that this was an important factor in the fall in consumption during the Great Recession (Heathcote and Perri, 2018). In existing models of this mechanism households precisely co-ordinate their labour market expectations on a particular equilibrium, as is common in models with multiple equilibria (Morris and Shin, 2000).

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

There are two important assumptions that drive this result. First, we assume that the hiring rate out of unemployment, which is crucial for precautionary saving in models with frictional labour markets (e.g. Ravn and Sterk, 2020), is not observed directly by households. Instead, we assume that households can obtain signals about the hiring rate, but that this information processing is costly. Households therefore choose to process somewhat noisy signals before deciding on their consumption, following the literature on rational inattention (Sims, 2003). Importantly, we allow households to choose the structure of their signals, not just how much information those signals contain. The hiring rate is naturally bounded between 0 and 1, and in rational inattention problems with a bounded unknown variable the optimal signal has a discrete number of possible realisations even though the underlying variable is continuous (Matějka, 2016). This nonlinear signal structure is what drives the possibility of multiple steady states.

The information processing cost implies that households have limited information about *realisations* of the hiring rate. Our second key assumption is that households also have limited information about the *structure* of their environment: they do not precisely know the true equilibrium marginal distribution of the hiring rate. This

is related to the ‘internal rationality’ studied by Adam and Marcet (2011). To our knowledge, we are the first paper to examine this combination of rational inattention and imprecise prior beliefs. This second assumption is consistent with hiring rate expectations reported in the Survey of Consumer Expectations: expectations have a much greater variance than the underlying hiring rate, which suggests that households do not have a good understanding of the range of values usually observed for the hiring rate.

The information choices made by households in this environment imply labour market expectations which are also consistent with survey data along other key dimensions. Expectations in the model are heterogeneous and persistent. They are only weakly co-ordinated across households: when the hiring rate falls, households on average receive signals that indicate it has fallen, and so average hiring rate expectations fall. These signals, however, are very noisy, so at the same time as the average expectation is falling a little, some households are shifting their expectations up, and others are expecting a large collapse in the hiring rate.

We find support for all of these features of expectations in the Survey of Consumer Expectations. Even after controlling for a range of personal characteristics, there is a great deal of disagreement about the expected hiring rate each month, and expectations are highly persistent at the household level. We also find evidence that co-ordination of beliefs is weak: changes in average expectations account for just 0.1% of the variation in household belief changes, suggesting that households do not simultaneously agree on shifts in the hiring rate. That is, most changes in household beliefs are idiosyncratic, as in the model.

The combination of information processing costs and imprecise prior beliefs is central to our results. To see why, consider an economy at a non-stochastic steady state. The true distribution of the hiring rate contains a single point. If households have precise knowledge of this distribution, then they know the hiring rate with certainty. The model with rational inattention is trivially identical to the model with full information, and the only possible steady state is the unique steady state from the full information model.

More generally, if small shocks imply that the hiring rate has a non-degenerate distribution that remains tightly concentrated around some steady-state level, the actions of a fully-informed household will also be concentrated in a tight range. If an inattentive household knows this distribution there is no need for them to consider actions far away from that full-information range. The actions of inattentive households therefore remain very close to the actions that would be taken under full information, so there is little scope for multiplicity. When households do not

know the true distribution of the hiring rate, however, their actions may deviate substantially from those that would be made under full information, and in that case it is possible for there to be multiple steady states, even when the model with full information has a unique steady state.

Given this logic, it is important for our results that beliefs about the hiring rate distribution remain imprecise over time. In the model we find that this is indeed the case, even when households can use the information they processed in previous periods to update their prior beliefs over time, because each period households process information until its marginal benefit equals its marginal cost. As long as that occurs before prior beliefs converge to the true distribution of the hiring rate, households will never choose to process the information required to learn that true distribution. This means that multiple steady states can survive even when beliefs are updated over time. When the information cost is such that there are three steady states, we find that the outer steady states (the unemployment and employment traps) are locally stable, while the middle steady state is unstable, with respect to ‘mutation’ shocks which alter the distribution of prior beliefs in the population. Large mutations can cause the economy to transition between steady states.

Section 2 places this work in the context of the literature. In Section 3 we develop our results within a simple model. In Section 4 we introduce our mechanism into a version of the HANK model from Ravn and Sterk (2020), in which the future hiring rate is very important in household decisions. We show that the combination of rational inattention and imprecise prior beliefs about the hiring rate generates three possible steady states: an unemployment trap with a low hiring rate and mild deflation, a middle steady state with higher employment and moderate inflation, and an ‘employment trap’ with very high employment and high inflation. In Section 5 we show that several key features of our model are found in survey data on hiring rate expectations. Section 6 concludes.

2 Related literature

This paper relates to several strands of existing literature. Firstly, there is a vast literature studying fluctuations and traps driven by self-fulfilling expectations (see Cooper and John (1988) for a review of the early literature). Specifically, in our model changes in labour market expectations affect precautionary saving decisions, which in turn affect aggregate demand and so the labour market. Challe and Ragot (2016) show that a tractable model featuring this mechanism fits US data on ag-

gregate 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. (2018) 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 (2020), who show that self-fulfilling labour market expectations can lead to the existence of multiple steady states: an economy can get stuck in a bad steady state where pessimistic beliefs persist indefinitely. We contribute to this literature by showing that these unemployment traps can be generated purely through the optimal information choices of rationally inattentive households, even in a model which is linear under full information. Moreover, our model does not require the co-ordination of beliefs which Morris and Shin (2000) argue is often necessary in models with multiple equilibria.

We also contribute to the literature on rational inattention (RI). Most existing models with RI over continuous variables have agents with quadratic objective functions collecting costly information about a random variable with a known Gaussian distribution.¹ This has proved useful in understanding consumption patterns (Sims, 2003; Luo, 2008), price stickiness (Maćkowiak and Wiederholt, 2009), business cycles (Maćkowiak and Wiederholt, 2015), and other macroeconomic phenomena. Matějka (2016) and Jung et al. (2019) show, however, that assuming a bounded prior belief leads to very different results to the quadratic-Gaussian formulation. Specifically, they show that the optimal decision rule of an agent facing a rational inattention problem with a bounded prior entails the agent choosing to limit themselves to a discrete number of options, even when the optimal choice under perfect information is continuous. As the probability of finding a job is naturally bounded by 0 and 1, our model displays these features. This paper is therefore a response to the call in Sims (2006) to explore the implications of RI away from the quadratic-Gaussian case. To our knowledge we are the first to incorporate rationally inattentive households with bounded prior beliefs into a general equilibrium setting.²

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

¹This is convenient as the optimal posterior belief about the shock, after processing information, is also Gaussian. See Maćkowiak et al. (2020) for a review of the RI literature.

²Stevens (2019) explores a model in which the opportunity for firms to pay a fixed cost to exactly reveal the state of the economy places a bound on the range of shock realisations under which a given signal structure is expected to be in operation, which leads to the same kind of discrete decision rules for her firms that we obtain for households.

this to explain movements in house prices and the current account. We extend this literature by showing that the interaction of internal rationality and information processing costs creates multiplicity, where neither assumption generates this by itself. We believe that we are the first to combine these two information restrictions. We deviate from existing literature on internal rationality in that we do not assume that agent beliefs are close to the true equilibrium distribution of the state variable. The typical logic for ‘near-rationality’ is that agents learn over time, so would discard any beliefs which are far from 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 (2019) show that inflation expectations are heterogeneous across households; in Section 5 we document that the same is true for hiring rate expectations. The theoretical implications of heterogeneity in inflation expectations have been studied by Andrade et al. (2019) and Wiederholt (2015), among others. In contrast to this literature, we study heterogeneous labour market expectations.

3 Static model

This section illustrates our mechanism in a simple static model. There are two sub-periods in this setup, which we refer to as morning and afternoon. The consumer problem is related to Ravn and Sterk (2020) and to our HANK model in Section 4, but is substantially simplified by the static nature of the problem. The firm side is also kept very simple, to illustrate the key forces driving the unemployment trap. In Section 3.8 we construct a dynamic model by repeating this static model, allowing households to update their prior beliefs over time using information processed in previous periods. This demonstrates that the unemployment trap does not disappear even when households update their prior beliefs over time.

3.1 Households

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

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

$$c_{mi} + \frac{b_{mi}}{R} = 1 \quad (1)$$

$$c_{ai}^e = w + b_{mi} \quad (2)$$

$$c_{ai}^u = \theta + b_{mi} \quad (3)$$

Combining these we have:

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

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

Households have quadratic utility over consumption in each sub-period: $U(c_{ti}) = -\frac{1}{2}(c_{ti} - c_0)^2$, where c_0 is a constant. They discount afternoon consumption with a discount factor β , so they choose morning consumption c_{mi} to maximise:

$$\mathbb{E}_{mi} V(c_{mi}, \eta) = \mathbb{E}_{mi} \left(-\frac{1}{2}(c_{mi} - c_0)^2 - \frac{\beta}{2}\omega(1-\eta)(c_{ai}^u - c_0)^2 - \frac{\beta}{2}(1-\omega(1-\eta))(c_{ai}^e - c_0)^2 \right) \quad (6)$$

Where c_{ai}^e and c_{ai}^u are given by equations (4) and (5). $V(c_{mi}, \eta)$ denotes the expected present value of choosing morning consumption c_{mi} when the hiring rate is η . The hiring rate is the only variable about which the household may be uncertain.

A fully-informed household observes the hiring rate, and chooses morning consumption according to:

$$c_m^{FI} = \frac{\beta R}{1 + \beta R^2} (R + w - \omega(1 - \eta)(w - \theta)) + \frac{c_0(1 - \beta R)}{1 + \beta R^2} \quad (7)$$

Full-information morning consumption is therefore an increasing linear function of the hiring rate, due to a simple consumption smoothing motive.³

³When η rises the probability of being employed rises, which increases expected afternoon consumption, and so in turn increases optimal morning consumption. Throughout the paper we refer to the saving households engage in to insure themselves against future unemployment as ‘precautionary saving’, but we do not mean by this that the certainty equivalent level of afternoon consumption is below the expected level of c_a , which would require a utility function with a positive third derivative. Rather, we mean that households choose to save in this model as a precaution against future unemployment.

3.2 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 that doing so is costly. This cost is increasing in the informativeness of the signal chosen. This is formalised in equation (11) below. As well as the amount of information in the signal (which we will denote κ), the agent must also choose how this information is to be structured: they could choose a signal which is very accurate in some ranges of η but not in others, for example.

The payoff function being maximised is as in equation (6). The agent views the hiring rate as exogenous.⁴ 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 realisations to consumption choices.

This household problem is closely related to the firm profit maximisation problem in Matějka (2016). We therefore proceed by following the steps used to solve Matějka’s firm problem. First, we simplify the household problem by noting that the quadratic value function in equation (6) implies a one-to-one mapping between the expected hiring rate and optimal morning consumption. A household will never choose a signal structure that has two distinct possible realisations that imply the same expected hiring rate, because distinguishing between the two realisations is a waste of information processing. There will therefore be a one-to-one mapping between signal realisations and the optimal morning consumption choice. We can therefore leave the signal choice in the background of the problem, and instead study the optimal decision rule linking morning consumption to the hiring rate, subject to the information costs of implementing such a rule.

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

⁴The hiring rate will in fact be endogenous to aggregate agent choices, but the agent does not take this into account. This is explored in more detail in Section 3.7.

The problem of household i is therefore:

$$f_i = \arg \max_{\hat{f}_i} \mathbb{E}_{c_{mi}}[V(c_{mi}, \eta)] - \psi \kappa(\hat{f}_i, g_i)$$

$$= \arg \max_{\hat{f}_i} \int_{\eta} \int_{c_{mi}} V(c_{mi}, \eta) \hat{f}_i(\eta, c_{mi}) d\eta dc_{mi} - \psi \kappa(\hat{f}_i, g_i) \quad (8)$$

subject to

$$\int_{c_{mi}} \hat{f}_i(\eta, c_{mi}) dc_{mi} = g_i(\eta) \quad \forall \eta \quad (9)$$

$$\hat{f}_i(\eta, c_{mi}) \geq 0 \quad \forall \eta, c_{mi} \quad (10)$$

$$\kappa(\hat{f}_i, g_i) = H[g_i(\eta)] - \mathbb{E}_{c_{mi}} H[\hat{f}_i(\eta|c_{mi})] \quad (11)$$

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

$$H[g_i(\eta)] = - \int g_i(\eta) \log g_i(\eta) d\eta \quad (12)$$

The first constraint (9) ensures that the marginal distribution of the hiring rate η obtained from the optimal joint pdf is consistent with $g_i(\eta)$, household i 's prior belief about the distribution of the hiring rate.⁵

The second constraint (10) is that the solution must be positive everywhere, as required for the decision rule to be a joint pdf.

The final constraint (11) is the information processing constraint. Entropy $H[\cdot]$ is a measure of the dispersion of a distribution. The first term of constraint (11) is the entropy of the prior. The prior reflects the information held by the household about the distribution of the hiring rate before receiving any signals. We will assume for now that this prior is the same for all households and is uniformly distributed, so the prior is rather dispersed and entropy is high (this is relaxed in Section 3.8). The second term is the expected entropy of $f_i(\eta|c_{mi})$, the updated distribution over the hiring rate believed by the household after processing their signals.⁶ A precise posterior knowledge of η would give a very low posterior entropy, so the entropy difference from the prior would be large. Information costs in this model

⁵In Matějka (2016), this marginal distribution of the variable subject to rational inattention is the true distribution of that variable. This will not be the case here, as η will be determined endogenously in the model. Instead, the marginal distribution obtained by integrating the joint pdf over consumption should be interpreted here as the distribution of the hiring rate the household is expecting to see from their 'ignorance prior' (see Section 3.7).

⁶The information content of the signals is incorporated into the choice of c_{mi} , so the conditional distribution of η given the choice of c_{mi} gives the posterior belief about η . This is a consequence of the one-to-one mapping between signals and actions discussed above.

are proportional to this difference, that is how much the agent can learn from the signals. Note that with identical prior beliefs, each household now faces exactly the same problem, and so each household chooses the same decision rule f_i . We therefore drop the i subscripts from the decision rule and the prior belief g . Households will still choose heterogeneous values for morning consumption, because the optimal signals contain idiosyncratic noise.

Given a particular level of information processing κ , the household must decide how to allocate that information processing budget. They could, for example, ensure they make no mistakes at all when the hiring rate is above a certain threshold, but in doing so they must accept that they will make larger mistakes with higher probability when η is below that level. This means that the household faces a trade-off: for a particular κ they can distinguish between a several values of η which are close together, but that reduces the entropy of the posterior a great deal, so outside of that small range of η their posterior $f(\eta|c_{mi})$ must remain dispersed. When η is in that small range, the household making that decision will be very accurate in choosing optimal c_{mi} , but when η is outside of that range they will make large mistakes with a high probability. Alternatively, they can choose to allocate their information processing capacity to distinguishing between a small number of cases which are far apart. They are then never precise in predicting the hiring rate, but they make large mistakes less often. This is what drives the result in Matějka (2016) that the agent optimally restricts themselves to a small number of discrete levels of the choice variable, even though a continuous range of that variable is available, when the optimal κ is sufficiently small that the information constraint binds.

3.3 Rational inattention solution

The hiring rate is naturally bounded by 0 and 1. Assume that prior beliefs are uniform over the whole of this range: $g(\eta) \sim U[0, 1]$. The optimal decision rule is plotted in Figure 1 for the parameters specified in Appendix A. The marginal cost of information ($\psi = 0.002$) is sufficiently high that the optimal information strategy is to collect signals which are less than perfectly informative about η . The information processed at this cost is such that the household optimally restricts themselves to two levels of morning consumption, even though under perfect information the optimal morning consumption choice is continuous in the hiring rate. The logic behind this is discussed in Section 3.2 above, and in detail in Jung et al. (2019) and Matějka (2016). As the hiring rate increases (and so the optimal choice of morning consumption under full information increases), the probability a household chooses the higher level of morning consumption in this restricted menu increases.




Figure 1: Optimal decision rule for $\psi = 0.002$. Other parameters in Appendix A.

We will refer to the low and high levels of morning consumption chosen with positive probability as c_L and c_H respectively. The probability of choosing each level of morning consumption conditional on η is (proof in Matějka, 2016):

$$f(c_j|\eta) = \frac{\hat{V}(c_j, \eta)f(c_j)}{\hat{V}(c_j, \eta)f(c_j) + \hat{V}(c_{-j}, \eta)(1 - f(c_j))}, \quad j \in \{L, H\} \quad (13)$$

Where we have defined $\hat{V}(c_j, \eta) = \exp\left(\frac{V(c_j, \eta)}{\psi}\right)$, and $f(c_j) = \int_0^1 f(\eta, c_j)d\eta$ is the unconditional probability of choosing c_j .

In this simple model, quadratic utility and the uniform prior mean that there is no incentive to err on the side of over- or under-consumption. We find numerically that $f(c_j) = 0.5$, and c_L and c_H are an equal distance below and above the morning consumption a fully-informed household would choose when $\eta = 0.5$.⁷

3.4 Aggregate consumption

There is a unit mass of households making this decision. They all face the same labour market conditions, but we assume that the noise in their signals is idiosyncratic. Therefore for each level of the hiring rate some agents choose each of the morning consumption levels in the optimal menu which arises from the RI problem with uniform priors (equations (8) - (11)). The proportions on each level of morning

⁷This is unsurprising, since Proposition 3 in Matějka (2016) implies that agents will more frequently take actions associated with regions of η where $-\left(\frac{\partial c_{mi}^{FI}}{\partial \eta}\right)^2 \frac{\partial^2 V(c, \eta)}{\partial c^2}$ is relatively low. In our model this function is constant at $\frac{(\beta R \omega(w - \theta))^2}{1 + \beta R^2}$, and so there is no such bias.

consumption are determined by the probabilities in the optimal joint pdf obtained as the decision rule from the household problem.

Therefore for each level of η we obtain aggregate morning consumption \bar{c} using:

$$\bar{c}(\eta) = \int_{-\infty}^{\infty} c_{mi} f(c_{mi}|\eta) dc_{mi} \quad (14)$$

Households choose morning consumption to maximise $E_{mi}V(c_{mi}, \eta)$ given their information set. With no information processing, that information set consists of their prior belief only.⁸ With a uniform prior $E_{mi}\eta = 0.5$ for all households, regardless of the realised value of η . Households therefore do not change their morning consumption choice as the hiring rate varies. As the optimal information processing capacity κ increases (as ψ decreases), some information about η is processed, so households begin to choose different levels of c_{mi} for different underlying η . For low values of κ , households optimally restrict themselves to two values of morning consumption. Importantly, the aggregate morning consumption function has a wave-like shape around its full information equivalent. This is shown in Figure 2, which plots the $\bar{c}(\eta)$ implied by the decision rule in Figure 1, alongside the corresponding aggregate morning consumption functions if households have no information and full information about η .




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

⁸We make the standard assumption in the RI literature that households cannot process negative amounts of information - they cannot choose to forget information in their prior belief. That means that if ψ exceeds the marginal utility of information when the information set only contains the prior belief, households hit this no-forgetting constraint and process no information ($\kappa = 0$).

Consider the case where information processing is constrained but non-zero (the blue curve). This meets the full information morning consumption function (in red) at $\bar{c} = 1.06, \eta = 0.5$. However, at this hiring rate in the full information model, every household consumes the same amount. In contrast, in the rational inattention model, half of the households get a signal that the hiring rate is ‘high’ and choose the high level of $c_{mi} = c_H = 1.105$. The other half receive a signal that η is ‘low’, and so consume the lower level $c_{mi} = c_L = 1.013$.

The flatter sections of the aggregate morning 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 morning consumption in their menus. In Figure 1 above, it can be seen that this is the case for extreme high and low values of η , and $\bar{c}(\eta)$ is flat in these regions accordingly. In contrast, as η moves from 0.4 to 0.6, large numbers of households switch from choosing the low level of morning consumption to the high level, and this corresponds to the steep section of the corresponding aggregate morning consumption function in Figure 2.

The shape of the aggregate morning 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 morning consumption in Figure 1 fell linearly as η increased $\bar{c}(\eta)$ would be linear. In fact, the distribution of $c_{mi}|\eta$ for the values of c_{mi} in the optimal menu is logistic in shape, which is what gives rise to the wave-like shape of the aggregate response curve.⁹ 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.

As shown in Matějka (2016), as information processing κ rises further, more choices of c_{mi} are introduced into the optimal menu. As this occurs the aggregate response of morning consumption to the hiring rate approaches the perfect information first order condition. For ease of exposition, the graphs in this paper are all drawn for information costs that imply households choose an optimal menu with two levels of consumption, but this is not important for the results, as conditional choice probabilities remain logistic. An example with a lower information cost in this static model is shown in Appendix B.

3.5 Firms

Usually, rational inattention models specify that agents collect costly information about exogenous variables, often shock processes. In contrast, we assume that the

⁹Matějka and McKay (2015) study in detail the links between RI and the logit model.

hiring rate is determined in equilibrium by firm hiring decisions.¹⁰ For the purposes of this simple model, it will be sufficient to say that firms hire more workers when aggregate morning demand is high, so:

$$\eta = H(\bar{c}), \text{ with } H'(\bar{c}) > 0 \quad (15)$$

The focus of this model is household information choices, and the aggregate morning 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 C microfoundations such a process for the hiring rate using a model with working capital. In Section 4 the firm side of the model is standard, as in Ravn and Sterk (2020).

3.6 Equilibrium

It is important here that the hiring rate η is not necessarily uniformly distributed like household prior beliefs. In Matějka (2016) and many other models in the rational inattention literature, agents' prior beliefs about the distribution of the unknown variable are correct. We extend the RI literature by considering prior beliefs which do not precisely match the true equilibrium distribution of the relevant variables. This may be plausible 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 5 suggests that for the hiring rate, households do not know the true equilibrium distribution.

In particular, we begin by assuming that households have a uniform 'ignorance' prior, which would be justified if the households do not understand how their decisions (and those of other agents) affect the hiring rate.¹¹ This is discussed in Section 3.7. We study equilibrium in the static model with these uniform priors. In Section 3.8 we take this as a starting point and repeat the static problem many times, allowing households to update their prior beliefs each period. We show that the multiplicity does not disappear if agents update their priors away from the uniform starting point. Since the household (14) and firm (15) conditions are functions of the hiring rate and morning consumption only, afternoon consumption is irrelevant

¹⁰This distinction is irrelevant in a standard model where agents perfectly understand how endogenous variables are determined, 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.

¹¹We require a prior belief which is significantly dispersed, but it does not need to necessarily be uniform for these results. The prior belief must be bounded, which is ensured in this case as the hiring rate is between 0 and 1 by definition.

for equilibrium. From this point, we therefore refer to morning consumption as consumption, and drop all m subscripts.

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

The graph below shows this equilibrium interaction. The blue and red aggregate consumption functions are as in Figure 2. The firm condition (15) is added in black. Under full information there is one equilibrium, but under RI there are three: an unemployment trap with low consumption, a middle equilibrium, and an employment trap with very a high consumption. The three equilibria are associated with low, medium, and high levels of the hiring rate respectively.




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

Consider first the middle equilibrium under rational inattention, at $\eta = 0.48$. As discussed in Section 3.4, at this hiring rate, close to the middle of the uniform prior belief, 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 can be seen when η falls from this central equilibrium. Under full information, all households respond the same way, by reducing their consumption linearly with the fall in η . With rational inattention, as discussed in Section 3.4, the response is heterogeneous and non-linear. As η falls initially, large numbers of households switch

from the high level of consumption c_H to the low level c_L , so aggregate consumption falls a great deal. But as η falls further, there is less change in the proportions on c_H and c_L . Aggregate consumption therefore declines more gradually, which leads to another equilibrium at $(\eta = 0.19, \bar{c} = 1.014)$ in which almost all households choose the low level of consumption. This is the unemployment trap. Similarly, there is a high employment steady state at $(\eta = 0.89, \bar{c} = 1.105)$ which is also not present under full information.

In effect, households are using their limited information processing capacity to decide if they face a ‘high’ or ‘low’ hiring rate. As η moves a little below 0.5, the majority of agents decide on ‘low’, and consume accordingly, whereas if they knew the hiring rate more precisely they would choose consumption based on only a ‘somewhat low’ η . As the hiring rate falls even further, the proportion on ‘low’ grows more slowly, and households do not decide on even lower consumption. They continue to believe that the hiring rate is ‘low’ even when it becomes ‘extremely low’. The aggregate consumption function therefore flattens out at very low values of the hiring rate, which is why there is an equilibrium with very little labour market activity.

3.7 Requirements for multiplicity

The multiplicity of equilibria is driven by the non-linearity in the aggregate consumption function, which arises endogenously from the optimal information processing decisions of households. To obtain this, we require that households process costly information about the hiring rate, with an ‘ignorance prior’. That is, households have only imprecise knowledge of the true equilibrium distribution of the hiring rate.

If information processing is not costly ($\psi = 0$), then households process information until they know η for certain. Equation (7) shows that the resulting full-information aggregate consumption function is linear in η , so there can only be one equilibrium. Similarly, if the cost of information is too high, households will process no information, and their consumption choices will not respond at all to the realised hiring rate. Again, the aggregate consumption function is linear and there is a unique equilibrium. Multiple equilibria of the kind we identify are therefore only possible if households process positive, but limited, amounts of information. In that case, households choose discrete nonlinear signals of the form studied above, the aggregate consumption function is nonlinear, and so multiple equilibria are possible.

To see the importance of the ignorance prior, consider the equilibria in Figure 3. 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 (the optimal value when $\eta = 0.5$) than it is in the figure, as they do not need to worry about the possibility of an extremely low η . This, in turn, would mean that the equilibrium value of η in the unemployment trap would be higher. Iterating this logic, household prior beliefs get more and more precise until the marginal benefit of the first unit of information processing falls below the marginal cost ψ . At that point households process no information, so the equilibrium must be unique.¹²

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

Our assumption of ‘ignorance priors’ builds on these existing rational inattention models by assuming that households do not perfectly understand the links from shocks to aggregate household choices to the hiring rate.¹³ This assumption is related to the removal of ‘external rationality’ in Adam and Marcet (2011). Like them, we believe that it is plausible that agents do not precisely know the true stochastic processes and mechanisms that determine the endogenous variables they face, especially when those mechanisms are complicated and involve the behaviour of many other agents. The survey data on hiring rate expectations in Section 5 supports this view: the variance of survey expectations (even after accounting for a variety of household characteristics) is an order of magnitude larger than the long run variance of the true hiring rate. It appears that, just as in our model, households do not know the appropriate range of values for the hiring rate.

However, even if households do not understand the mechanisms leading to a particular equilibrium distribution of the hiring rate, they might obtain accurate prior beliefs if they observed the hiring rate over many periods, and learned its equilib-

¹²The notion that information in priors substitutes for costly new information is explored further below and in Section 3.8. In principle, we believe it is possible that with certain calibrations and carefully-chosen processes for selecting between equilibria, the iteration described here would not imply beliefs converging to a single point. In that case multiple equilibria could exist without the ignorance prior. Study of when these equilibria exist and their properties is beyond the scope of this paper. The results in Sections 3.8 and 3.9 require that beliefs remain substantially dispersed even in a deterministic steady state, where the true distribution of η contains a single point.

¹³Note that with ignorance priors, higher-order beliefs cannot matter for household choices because households do not understand how the actions of others feed into the hiring rate.

rium distribution over time. This would imply ‘near-rational’ subjective beliefs, which are sufficiently close to the true equilibrium process that agents cannot distinguish between the two given the data they observe (as in Adam and Marcet, 2011). In Section 3.8 we show that this does not occur, and beliefs remain substantially dispersed even if beliefs can update over time, as long as households can change their information processing over time. This is because the updating of prior beliefs from the initial period (when priors are uniform) implies households process less information in subsequent periods. Households rely more on their priors and less on new signals, because their prior beliefs are more informative than in that initial period (Steiner et al. (2017) explore a similar mechanism in a discrete choice environment). Prior beliefs therefore do not approach the true distribution of the endogenous variable.¹⁴

The other important condition for multiplicity is that the firm hiring function $H(\bar{c})$ must be upward sloping. The aggregate consumption function under rational inattention is non-linear, but is always upward sloping. If it was downward sloping at any point, then it must be the case that an increase in the hiring rate leads 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 $dc^{FI}/d\eta > 0$. The rational inattention aggregate consumption function cannot therefore be downward sloping. This means that if $H'(\bar{c}) < 0$, there will be a unique equilibrium.

We therefore need a degree of strategic complementarity between households and firms for our multiplicity results. This is very plausible in models of precautionary saving based on labour market expectations: Ravn and Sterk (2020) 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 (2020), there is an unemployment trap steady state at a hiring rate of zero because the Phillips Curve is kinked at that point by the requirement that vacancy posting cannot be negative.¹⁵ Our model is

¹⁴The 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).

¹⁵In their Appendix A.1.2 they generate an unemployment trap steady state with a positive hiring rate through a different non-linearity that enters the household Euler equation when the assets held by households have a positive net supply. In Section 4 we focus on the case where the net supply of assets is zero, so this additional nonlinearity is not present.

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 any exogenously imposed non-linearity in another part of the model. The non-linearity arises endogenously from optimal household information choices. For this reason, in Section 4 we find an unemployment trap with low, but positive, labour market activity, whereas the unemployment trap in Ravn and Sterk has zero employment. We also find a high employment trap.

3.8 Dynamic solution

Our multiplicity result requires that households have ignorance priors. In a dynamic model where households process information about the hiring rate each period, they might update their prior beliefs using information processed in previous periods. Our multiple equilibria will only therefore survive in the long run if this belief updating does not lead to beliefs converging to the true distribution of the hiring rate over time. In this section, we repeat the static model from Sections 3.1-3.6, allowing households to learn about the distribution of the hiring rate over time. Choosing a belief updating process that ensures the model remains tractable, we show that prior beliefs do not converge to the true distribution of η , even in the long run. Multiple steady states persist even when households update their priors over time.

Define a period as an iteration of the static model, composed of a morning and afternoon. As in Section 3.1, all households are employed each morning, and some lose their job at the end of the morning. Firms then hire new workers for the afternoon in a frictional labour market. Households start from the uniform prior belief in period 1, but they can use information from one period to update their prior beliefs about the distribution of η for the next period. Prior beliefs will not therefore be uniform in period $t > 1$. As in the static model, households consume all of their income and savings each afternoon, so households do not build up wealth, and prior beliefs are the only link between periods. Again, afternoon consumption is irrelevant for equilibrium, so we use consumption to mean morning consumption.

Households with ignorance priors do not know the true process generating the hiring rate, so they cannot update their beliefs using the Bayesian updating rule derived from that process. Instead, we consider a simple rule for updating beliefs, which we view as a rule-of-thumb behaviour from households who know very little about the dynamic process for η . If household i has posterior belief $f(\eta_t|c_{i,t})$ after processing signals in period t , we will suppose that their prior belief before

information processing in period $t + 1$ is given by:

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

That is, 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 . This particular updating rule is useful because it implies the model remains tractable over many periods, as will be shown below.

A further implication of ignorance priors is that households cannot accurately compute the future value of information processed in the current period. We therefore follow Kreps (1997) in assuming households maximise 'anticipated utility'. That is, households in period t do not take into account the informational value of their choices for future periods. This is a standard way to avoid the intractability that comes from optimisation problems where current choices generate information which may be useful in future periods (see e.g. Cogley and Sargent, 2005).

With this assumption households act as if they face a series of unconnected static problems. Prior beliefs are the only connection between periods. In the initial period, households have a uniform prior so their problem is identical to the one in Sections 3.2 and 3.3. Households choose a signal with two possible realisations, which imply consumption choices of c_L and c_H . We obtain the corresponding posterior beliefs $f(\eta_1|c_{L,1})$ and $f(\eta_1|c_{H,1})$ by applying Bayes' rule to equation (13):

$$f(\eta_1|c_{j,1}) = \frac{2\hat{V}(c_{j,1}, \eta_1)}{\hat{V}(c_{j,1}, \eta_1) + \hat{V}(c_{-j,1}, \eta_1)}, \quad j \in \{L, H\} \quad (17)$$

Through the updating rule (16), each of these two possible posteriors translates into a prior belief in the second period. Prior beliefs differ from the uniform distribution of period 1 (as long as $\rho > 0$), and households will make different information decisions in period 2 than in period 1.

Recall that agents choose how much information to process (κ) such that the marginal benefits of more information equal the marginal cost ψ . Information in prior beliefs is a substitute for information from new signals. This means that 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, as they have incorporated some of the information processed in period 1. Households therefore process less information in period 2 than they did in period 1, and rely more on their priors to guide their decisions. Importantly, this implies that priors do not converge to the true distribution of η over time. Each

period households process enough information to return them to the point where the marginal benefit of information equals ψ . The amount of information in posteriors, and so priors, does not therefore increase over time. Since prior beliefs have not collapsed to the truth in period 2, they will not do so in any future period.¹⁶

The set of posterior beliefs held by households at the end of period 2 is in fact *identical* to the set held in period 1 when the updating rule is (16). That is, whether a household begins period 2 with prior $g(\eta_2|c_{L,1})$ or $g(\eta_2|c_{H,1})$, they choose a signal with the same two possible realisations as in period 1. Those realisations imply households choose from the same menu of consumption choices as they did in period 1, and the posterior belief formed after choosing c_j in period 2 is the same as it was in period 1. We prove this in Appendix D.1.

Since the set of posteriors is the same in period 2 as in period 1, the set of prior beliefs in period 3 is the same as it was in period 2. Households therefore make the same information choices in period 3 as they did in period 2, and so again the menu of consumption choices and set of posterior beliefs remains the same. Iterating this argument for each subsequent period implies that households continue to be split between these two priors and posteriors. The prior belief $g(\eta_{t+1}|c_{j,t})$ is therefore independent of t and of signal realisations before period t .

The two priors for the calibration used to draw Figures 1-3, and $\rho = 0.9$, are plotted in Figure 4. Households receiving a signal in period t that the hiring rate is low form the prior belief $g(\eta_{t+1}|c_{L,t})$ (plotted in blue). This puts more weight on low values of η than the initial uniform prior because it has been formed using information from the period t signal. For the same reason, the prior of a household who chose c_H in period t (in red) has the most weight on high values of η . We refer to households holding the low prior $g(\eta_{t+1}|c_{L,t})$ as *pessimists*, and those holding the high prior $g(\eta_{t+1}|c_{H,t})$ as *optimists*.

¹⁶This insight is not specific to the belief updating rule in (16). A different rule would imply a different prior $g(\eta_2|c_{j,1})$, but as long as that prior has a greater entropy than the true distribution of η then households will never pay for the information required to learn that true distribution. The updating rule in (16) and the uniform initial prior are however required for the results we use to obtain analytic results for the dynamic model below.




Figure 4: Prior beliefs $g(\eta_{t+1}|c_{L,t})$ (blue) and $g(\eta_{t+1}|c_{H,t})$ (red) in period $t + 1 > 1$ with $\psi = 0.002$ and $\rho = 0.9$. Other parameters in Appendix A.

Although optimists and pessimists have the same set of possible posterior beliefs, they do not have the same conditional choice probabilities, because the marginal probabilities of each consumption level differ across priors. Specifically, we show in Appendix D.1 that:

$$f(c_{j,t+1}|c_{j,t}) = 0.5(1 + \rho) > 0.5(1 - \rho) = f(c_{j,t+1}|c_{-j,t}) \quad (18)$$

A household is therefore more likely to choose consumption c_j if they also chose it in the previous period. Intuitively, a household chooses c_L when the signal they receive implies the hiring rate is low. In the next period, that household has pessimistic priors, with a large weight on a low hiring rate, and so they are more likely to choose a low consumption again. Choices and expectations are therefore persistent at the individual level.

Since the prior belief in period $t + 1$ is determined by the posterior (and so consumption choice) in period t , the conditional choice probabilities for consumption are also the probabilities of a household entering period $t + 1$ with the prior belief $g(\eta_{t+1}|c_{j,t})$, conditional on η_t and the prior belief they held at the start of period t . Households therefore move between being pessimists and optimists according to the transition matrix in Table 1, where each cell (j, k) is $f(c_{k,t}|c_{j,t-1}, \eta_t)$, derived using (13) and the marginal probabilities from (18).

Table 1: Prior belief transition matrix

| | Pessimist _{t+1} | Optimist _{t+1} |
|------------------------|--|--|
| Pessimist _t | $\frac{(1+\rho)\hat{V}(c_L, \eta_t)}{(1+\rho)\hat{V}(c_L, \eta_t) + (1-\rho)\hat{V}(c_H, \eta_t)}$ | $\frac{(1-\rho)\hat{V}(c_H, \eta_t)}{(1+\rho)\hat{V}(c_L, \eta_t) + (1-\rho)\hat{V}(c_H, \eta_t)}$ |
| Optimist _t | $\frac{(1-\rho)\hat{V}(c_L, \eta_t)}{(1+\rho)\hat{V}(c_H, \eta_t) + (1-\rho)\hat{V}(c_L, \eta_t)}$ | $\frac{(1+\rho)\hat{V}(c_H, \eta_t)}{(1+\rho)\hat{V}(c_H, \eta_t) + (1-\rho)\hat{V}(c_L, \eta_t)}$ |

Let \mathbf{P}_t be the proportion of households who hold pessimistic priors at the start of period t . Using the definition in (14), the aggregate consumption function can be written as:

$$\begin{aligned} \bar{c}_t = & c_L(\mathbf{P}_t f(c_{L,t}|\eta_t, c_{L,t-1}) + (1 - \mathbf{P}_t)f(c_{L,t}|\eta_t, c_{H,t-1})) \\ & + c_H(\mathbf{P}_t f(c_{H,t}|\eta_t, c_{L,t-1}) + (1 - \mathbf{P}_t)f(c_{H,t}|\eta_t, c_{H,t-1})) \end{aligned} \quad (19)$$

Next we define the stationary distribution of beliefs for each η as the distribution of households between optimists and pessimists such that if the hiring rate remains constant at η , the belief distribution will also remain constant. Formally:

Definition. *The stationary distribution of beliefs associated with a hiring rate of η is a proportion of pessimists $\mathcal{P}(\eta)$ such that:*

$$\mathcal{P}(\eta) = \mathcal{P}(\eta)f(c_{L,t}|\eta, c_{L,t-1}) + (1 - \mathcal{P}(\eta))f(c_{L,t}|\eta, c_{H,t-1}) \quad (20)$$

We use these definitions to characterise the steady states of the model.

Definition. *A steady state is a hiring rate η , an aggregate consumption \bar{c} , and a distribution of beliefs \mathbf{P} such that:*

1. *The belief distribution is stationary given η : $\mathbf{P} = \mathcal{P}(\eta)$ satisfies (20).*
2. *Aggregate consumption is consistent with the stationary belief distribution, given η : \bar{c} satisfies (19).*
3. *The hiring rate and aggregate consumption are consistent with the firm condition: η and \bar{c} satisfy (15).*

To visualise the steady states of the model, it is helpful to define a steady state aggregate consumption function $\bar{c}^{ss}(\eta)$, which gives the aggregate consumption at hiring rate η when $\mathbf{P} = \mathcal{P}(\eta)$. The points along this curve therefore give the combinations of η, \bar{c} that satisfy requirements 1 and 2 in the definition of steady state.

In the stationary distribution, if the hiring rate remains constant then the proportion of households with pessimistic priors is equal to the proportion choosing c_L

each period, and so aggregate consumption with beliefs in their stationary distribution is:

$$\bar{c}^{ss}(\eta) = c_L \mathcal{P}(\eta) + c_H(1 - \mathcal{P}(\eta)) \quad (21)$$

Figure 5 plots this steady state aggregate consumption function (blue solid) and the firm condition (black dashed). Steady states occur where the two curves meet.




Figure 5: Steady state consumption function and firm condition. $\psi = 0.002$, $\rho = 0.9$, other parameters in Appendix A.

If beliefs did collapse to the truth, the steady state aggregate consumption function would coincide with the linear full information consumption function in Figure 2.¹⁷ In fact we retain the logistic shape of the aggregate response curve from the static problem, so this model generates multiple steady states. As with the static model, the nonlinearity in the consumption function comes from nonlinearities in the signals chosen by households. The probability that a household receives a low signal realisation does not vary strongly with η when it is close to 0 or 1, and so the stationary belief distribution is unresponsive to η in those ranges, implying a shallow aggregate consumption function. When η is close to 0.5, the probability a household receives a low signal realisation varies more with η , and so therefore does the aggregate consumption function.

With the parameters used to draw Figure 5, there is a high employment trap steady state with $\eta = 0.89$, a middle steady state with $\eta = 0.49$, and an unemployment trap steady state with $\eta = 0.18$. Table 2 gives the transition probabilities

¹⁷The true distribution of the hiring rate in deterministic steady state is a single point. If households know this, with no further information processing they choose the optimum consumption for that hiring rate.

between the two prior beliefs at each of the steady states (abbreviating Pessimist and Optimist as P and O respectively).

Table 2: Prior belief transition matrices

| (a) Unemployment trap | | | (b) Middle steady state | | | (c) Employment trap | | |
|-----------------------|-----------|-----------|-------------------------|-----------|-----------|---------------------|-----------|-----------|
| | P_{t+1} | O_{t+1} | | P_{t+1} | O_{t+1} | | P_{t+1} | O_{t+1} |
| P_t | 0.9995 | 0.0005 | P_t | 0.9514 | 0.0486 | P_t | 0.0619 | 0.9381 |
| O_t | 0.8919 | 0.1081 | O_t | 0.0565 | 0.9435 | O_t | 0.0002 | 0.9998 |

In the middle steady state, both optimistic and pessimistic beliefs are very persistent, so there is a great deal of belief heterogeneity in steady state. Each period 54% of households hold pessimistic prior beliefs and 46% hold optimistic beliefs, and individual households occasionally switch from one type to the other. In the unemployment trap, the hiring rate is so low that even households with optimistic prior beliefs are very likely to receive a signal that η is low, and so to switch to pessimistic beliefs. There is therefore much less heterogeneity in prior beliefs in the unemployment trap steady state: 99.94% of households hold pessimistic beliefs each period, though even here there is a small amount of churn in beliefs. Similarly, in the high employment steady state 99.98% of households hold optimistic beliefs.

3.9 Stability under belief mutations

After finding multiple steady states in this model, a natural follow-up question is how the system behaves when shocked away from those steady states. We analyse the stability of each steady state to shocks to the distribution of prior beliefs in the population, and find that the unemployment trap and high employment trap steady states are locally stable, while the middle steady state is unstable.

In this model the steady states are defined by the distribution of households among the two prior beliefs. In the spirit of Kandori et al. (1993) we define a *mutation* as a perturbation of the proportions of households holding pessimistic and optimistic priors, and we study the response of the endogenous model variables to such a shock.¹⁸ Specifically, suppose that at the beginning of each period a shock arrives that causes some pessimistic households to switch to holding optimistic priors, and some optimists to switch to pessimistic priors, before the households

¹⁸Kandori et al. (1993) study a 2x2 game with multiple Nash equilibria. They consider a large population of players meeting at random, and perturb the proportion of players choosing each strategy. While our setting differs from this, we similarly consider a large population of players (households) with two possible ‘types’ - prior beliefs instead of the strategies in Kandori et al. (1993).

choose their information strategy for the period. The mutation is therefore an exogenous shock to \mathbf{P}_t , the proportion of households who hold pessimistic prior beliefs before processing signals in period t .

One interpretation of these mutations that mirrors the intuition in Kandori et al. (1993) is that each period a fraction of households leave the economy, to be replaced by new households who have a different probability of being optimistic or pessimistic than the households they replaced. Alternatively, we could think that at the end of each period households start to doubt the signals that led them to their consumption choice. When they come to form their priors at the start of the next period, some decide they must have been wrong in the previous period and act as if they had received the opposite signal realisation.

The transition probabilities in Table 1 show that, absent mutation shocks, the proportion of pessimists in period $t+1$ is a function of the proportion in period t and the hiring rate. The equilibrium hiring rate, in turn, is a function of the distribution of beliefs, so we can write the mapping from \mathbf{P}_t to \mathbf{P}_{t+1} as:

$$\mathbf{P}_{t+1} = \phi(\mathbf{P}_t, \eta_t(\mathbf{P}_t)) \quad (22)$$

If a mutation shocks \mathbf{P}_t away from a steady state, then the effect on the proportion of pessimists in the next period is:¹⁹

$$\frac{d\mathbf{P}_{t+1}}{d\mathbf{P}_t} = \phi_{\mathbf{P}}(\mathbf{P}_t, \eta_t(\mathbf{P}_t)) + \phi_{\eta}(\mathbf{P}_t, \eta_t(\mathbf{P}_t)) \frac{d\eta_t(\mathbf{P}_t)}{d\mathbf{P}_t} \quad (23)$$

A steady state is locally unstable with respect to mutations if $\frac{d\mathbf{P}_{t+1}}{d\mathbf{P}_t}$ is greater than 1 in absolute value at that steady state.

We derive $\frac{d\mathbf{P}_{t+1}}{d\mathbf{P}_t}$ in Appendix D.2. First, we show that $|\phi_{\mathbf{P}}(\mathbf{P}_t, \eta_t(\mathbf{P}_t))| < 1$. That is, if the hiring rate does not respond to the mutation (if $\frac{d\eta_t(\mathbf{P}_t)}{d\mathbf{P}_t} = 0$), then any belief distribution is locally stable.

However, in general the hiring rate will respond to mutations. As the hiring rate rises, the probability of households choosing c_L falls, and so $\phi_{\eta}(\mathbf{P}_t, \eta_t(\mathbf{P}_t)) < 0$. A belief distribution will therefore be unstable if the equilibrium hiring rate responds sufficiently strongly to the mutation. Intuitively, if an increase in the proportion of pessimists causes a fall in the hiring rate, this increases the probability of all households choosing c_L , so increases the probability of them holding pessimistic

¹⁹For many values of \mathbf{P}_t there will be multiple hiring rates consistent with equilibrium. In this case we select the equilibrium in which the hiring rate $\eta_t(\mathbf{P}_t)$ is closest to its previous period value $\eta_{t-1}(\mathbf{P}_{t-1})$, following the inertia assumption in Kandori et al. (1993). This ensures that the hiring rate is differentiable, outside of a knife-edge case discussed in Appendix D.3.

priors in the next period. If the hiring rate response is sufficiently strong there will be even more pessimists in the following period and the initial belief distribution is unstable.

We find that at the outer steady states (the unemployment trap and the employment trap), the hiring rate response to mutations is sufficiently weak that the steady states are locally stable. In contrast, at the middle steady state the hiring rate responds strongly to mutations. The middle steady state is therefore locally unstable.²⁰

This analysis concerns small mutations near the steady state. In Appendix D.3 we consider large mutations, and show that a sufficiently large increase in the proportion of pessimistic households can cause the model to switch from the employment trap steady state to the unemployment trap. Similarly, a large increase in the proportion of households with optimistic priors can lift an economy out of the unemployment trap, leading to a permanently higher hiring rate and aggregate consumption.

4 HANK model

Here we study the HANK model in Ravn and Sterk (2020) (RS) with the addition of rational inattention to the hiring rate and prior beliefs which do not match the true equilibrium hiring rate distribution, to show that our results continue to hold in richer models than that studied above. The RS model is particularly useful because it remains tractable despite featuring the uninsurable labour market risk which is necessary to generate a precautionary savings motive.

The tractability is achieved in RS 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 et al., 2018). Employed asset-poor households are the

²⁰In Appendix D.2 we show that the middle steady state is always unstable as long as an increase in the proportion of pessimists lowers the hiring rate ($\frac{d\eta_t(\mathbf{P}_t)}{d\mathbf{P}_t} < 0$), which is true of most calibrations, including the one used to draw the figures in this section. It is however possible to construct calibrations in which $\frac{d\eta_t(\mathbf{P}_t)}{d\mathbf{P}_t} > 0$ at the middle steady state. The conditions for instability, though similar, are not exactly the same in this case. We detail them in Appendix D.2.

critical households as they remain on their Euler equation, so they price the bond in equilibrium.

We take this model as our starting point, but make a number of changes to adapt it for our analysis. First, we remove the constraint that vacancy posting cannot be negative. In RS, this constraint places a kink in the Phillips Curve at a hiring rate of zero, which gives rise to an unemployment trap steady state with zero hiring. Without this, the model has a unique steady state under full information.²¹ We show that rational inattention generates multiple steady states with positive hiring in this environment.

In addition, we only impose the borrowing constraint on unemployed households, and we assume that real wages are fixed rather than determined by Nash bargaining. With full information, allowing employed households to borrow makes no difference, because if one household wants to borrow at a given interest rate they all do, so there is no-one to lend to them and the asset market does not clear. Employed households still therefore hold zero assets in equilibrium. The assumption will, however, matter under rational inattention, because households will have heterogeneous expectations. Some employed households will accumulate wealth by lending to others. Assuming fixed real wages ensures that the resulting wealth heterogeneity does not imply wage heterogeneity through heterogeneity in the value of unemployment, and so keeps the analysis tractable.

4.1 Households

All households h choose consumption to maximise expected discounted utility:

$$U_t^h = \mathbb{E}_t \sum_{s=t}^{\infty} \beta^s \left(\frac{c_{hs}^{1-\mu} - 1}{1-\mu} - \zeta n_{hs} \right) \quad (24)$$

Here n_{hs} is an indicator equal to 1 if the household is employed in period s and 0 if they are unemployed, and ζ is the constant disutility of labour.

Households maximise this subject to a budget constraint and, if they are unemployed, a borrowing constraint:

$$P_s c_{hs} + \frac{b_{hs+1}}{R_s} = P_s (w n_{hs} + \vartheta (1 - n_{hs})) + b_{hs} \quad (25)$$

$$b_{hs+1} \geq 0 \quad \text{if } n_{hs} = 0 \quad (26)$$

²¹Note that vacancy posting will never be negative in steady states we find. Removing the explicit constraint simply removes the kink in the Phillips Curve which allows for a zero hiring steady state.

P_s and R_s are the price level and the gross nominal interest rate. Employed households can borrow or save by trading one-period bonds b_{hs+1} . Employed households receive the exogenous real wage w , while unemployed households receive the value of home production ϑ . In the calibration we study, the no-borrowing constraint is binding for all unemployed households.

Each period s an exogenous proportion ω of employed households lose their job, and an endogenous proportion η_s of unemployed households find a job.

We begin by analysing the problem when households have full information about the future path of all variables. Denoting the choices of employed and unemployed households as c_{es}, b_{es+1} and c_{us}, b_{us+1} , the Bellman equations for employed and unemployed households respectively are:

$$V^e(b_s, \eta^{s+1}) = \max_{c_{es}, b_{es+1}} \frac{c_{es}^{1-\mu} - 1}{1-\mu} - \zeta + \beta\omega(1 - \eta_{s+1})V^u(b_{es+1}, \eta^{s+2}) \\ + \beta(1 - \omega(1 - \eta_{s+1}))V^e(b_{es+1}, \eta^{s+2}) \quad \text{subject to } P_s c_{es} + \frac{b_{es+1}}{R_s} = P_s w + b_s \quad (27)$$

$$V^u(b_s, \eta^{s+1}) = \max_{c_{us}} \frac{c_{us}^{1-\mu} - 1}{1-\mu} + \beta(1 - \eta_{s+1})V^u(0, \eta^{s+2}) + \beta\eta_{s+1}V^e(0, \eta^{s+2}) \\ \text{subject to } P_s c_{us} = P_s \vartheta + b_s \quad (28)$$

Here η^s denotes the sequence $\{\eta_t\}_{t=s}^\infty$. $V^e(b_s, \eta^{s+1})$ and $V^u(b_s, \eta^{s+1})$ are the values of being employed and unemployed in period s with initial wealth b_s .

The consumption Euler equation for an employed household, substituting inflation $\Pi_{s+1} = P_{s+1}/P_s$, is as in RS:

$$c_{es}^{-\mu} = \beta \frac{R_s}{\Pi_{s+1}} \left(\omega(1 - \eta_{s+1})c_{u,s+1}^{-\mu} + (1 - \omega(1 - \eta_{s+1}))c_{e,s+1}^{-\mu} \right) \quad (29)$$

Employed households choose consumption in period s such that the marginal utility of consumption equals the expected marginal utility of future consumption over the two possible future labour market states, adjusted for discounting and the real interest rate. Unemployed households are always at their borrowing constraint, and so their problem never matters for equilibrium determination.

Interest rates are set using a simple interest rate rule:²²

$$R_s = \bar{R}\bar{\Pi}^{-\delta\pi}\Pi_s^{\delta\pi} \quad (30)$$

²²This is as in RS, abstracting from the zero lower bound on interest rates and setting the coefficient on the hiring rate in the interest rate rule to 0, as RS do in drawing their Figure 2.

To get the steady state Euler equation (**EE**) note that since equilibrium asset holdings are always zero $c_{es} = w$ and $c_{us} = \vartheta$. Using this and the interest rate rule (30) in equation (29) in steady state gives:

$$1 = \beta R^* \Pi^{\delta_\pi - 1} \left[\omega(1 - \eta) \left(\frac{\vartheta}{w} \right)^{-\mu} + 1 - \omega(1 - \eta) \right] \quad (31)$$

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

4.2 Firms

Firms in our model are identical to those in RS, except they do not face a non-negativity constraint on vacancies. Firms set prices and choose how many vacancies to post each period to maximise profits. There are quadratic price adjustment costs as in Rotemberg (1982), firms are monopolistic and households have CES preferences over the firms' output. The price setting and vacancy posting decisions lead to a Phillips Curve, which in steady state becomes (derivation in Appendix E.1):

$$\phi(1 - \beta)(\Pi - 1)\Pi = 1 - \gamma + \gamma(w + k\eta^{\frac{\alpha}{1-\alpha}}(1 - \beta(1 - \omega))) \quad (32)$$

Here ϕ measures the extent of price adjustment costs, γ is the elasticity of substitution between goods in the consumer's problem, k is the cost of posting a vacancy, ω is the (fixed) job separation rate, η is the hiring rate and q is the vacancy filling rate, equal to $\eta^{\frac{-\alpha}{1-\alpha}}$, where α is the elasticity of the Cobb-Douglas labour matching function with respect to job searchers.

Equation 32 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. λ_f is the Lagrange multiplier on the constraint that vacancies must be weakly positive, so is always zero in our version of the model. Steady-state vacancies will always be positive in our model.

4.3 Rational inattention

As in Section 3.2, we amend the household problem so that processing information about the next period hiring rate has marginal cost ψ .²³ Again, we start with uniform prior beliefs $g(\eta) \sim U(0, 1)$ for all households, then allow households to update their prior beliefs using a weighted average of a uniform distribution and

²³For simplicity we assume that they cannot process information about hiring rates further into the future.

the previous period's posterior, as in Section 3.8. The equations of the household problem therefore mirror the setup explained in Section 3.2.

Formally, the information problem of an employed household is:

$$\begin{aligned} f &= \arg \max_{\hat{f}} \mathbb{E}[\tilde{V}^e(b_s, \eta_{s+1}, b_{es+1}, g(\eta_{s+1}))] - \psi\kappa \\ &= \arg \max_{\hat{f}} \int_{\eta_{s+1}} \int_{b_{es+1}} \tilde{V}^e(b_s, \eta_{s+1}, b_{es+1}, g(\eta_{s+1})) \hat{f}(\eta_{s+1}, b_{es+1}) d\eta_{s+1} db_{es+1} - \psi\kappa \end{aligned} \quad (33)$$

subject to

$$\int_{b_{es+1}} \hat{f}(\eta_{s+1}, b_{es+1}) db_{es+1} = g(\eta_{s+1}) \quad \forall \eta_{s+1} \quad (34)$$

$$\hat{f}(\eta_{s+1}, b_{es+1}) \geq 0 \quad \forall \eta_{s+1}, b_{es+1} \quad (35)$$

$$\kappa = H[g(\eta_{s+1})] - \mathbb{E}_{b_{es+1}} H[\hat{f}(\eta_{s+1} | b_{es+1})] \quad (36)$$

Where:

$$\begin{aligned} \tilde{V}^e(b_s, \eta_{s+1}, b_{es+1}, g(\eta_{s+1})) &= \frac{(c_{es}(b_s, b_{es+1}))^{1-\mu} - 1}{1 - \mu} - \zeta \\ &\quad + \beta\omega(1 - \eta_{s+1})\mathbb{E}_{g(\eta_{s+1})} \tilde{V}^u(b_{es+1}, \eta_{s+2}, g(\eta_{s+1})) \\ &\quad + \beta(1 - \omega(1 - \eta_{s+1}))\mathbb{E}_{g(\eta_{s+1})} \tilde{V}^e(b_{es+1}, \eta_{s+2}, b_{es+2}, g(\eta_{s+1})) \end{aligned} \quad (37)$$

$$\begin{aligned} \tilde{V}^u(b_s, \eta_{s+1}, g(\eta_{s+1})) &= \frac{(c_{us}(b_s))^{1-\mu} - 1}{1 - \mu} + \beta(1 - \eta_{s+1})\mathbb{E}_{g(\eta_{s+1})} \tilde{V}^u(0, \eta_{s+2}, g(\eta_{s+1})) \\ &\quad + \beta\eta_{s+1}\mathbb{E}_{g(\eta_{s+1})} \tilde{V}^e(0, \eta_{s+2}, b_{es+2}, g(\eta_{s+1})) \end{aligned} \quad (38)$$

And:

$$c_{es}(b_s, b_{es+1}) = w + \frac{b_s}{P_s} - \frac{b_{es+1}}{P_s R_s}, \quad c_{us}(b_s) = \vartheta + \frac{b_s}{P_s} \quad (39)$$

Aside from substituting out for consumption using the budget constraints, the value functions differ from those of the full information problem (27 - 28) in two important ways. First, as households do not know the future path of the hiring rate, they are uncertain about the value functions they will face in the next period.²⁴ In forming these expected continuation values we assume that they use their period- s prior

²⁴We also assume that households don't observe future inflation or interest rates, but rather believe they will remain constant at current levels. For the steady state results presented here this is equivalent to assuming they know the future path of those variables as in the full information case, but it substantially simplifies the analysis of mutation shocks presented in Appendix E.4.

belief about the distribution of the hiring rate.²⁵ The value functions under RI therefore depend on the prior belief $g(\eta_{s+1})$. In keeping with the assumption of anticipated utility, the households believe that their prior in future periods will be equal to their current prior.

Second, the full information value functions assume that b_{es+1} is chosen optimally given η^{s+1} . However, as in Section 3.2, computing the expected utility of different information sets requires the payoffs from each possible combination of state of the world (η_{s+1}) and action (b_{es+1}). The value functions for the employed household's problem with RI are therefore defined over general choices of b_{es+1} . For the future expected value of being employed, households assume that they will choose b_{es+2} by solving the equivalent RI problem in the next period, with the same prior as they currently hold but new initial wealth. Taking the expectation of this value function using the prior $g(\eta_{s+1})$ also therefore implies taking expectations over the possible b_{es+2} choices the household might make, as well as over η_{s+2} .

Relaxing the no-borrowing constraint on the employed means that heterogeneity in the signals households receive implies a non-degenerate wealth distribution among the employed in steady state.²⁶ Some newly unemployed households will therefore have positive savings, but with our calibration they all hit their borrowing constraint when they become unemployed for all inflation rates and hiring rates consistent with steady state and equilibrium. Unemployed households therefore set $b_{us+1} = 0$ in all states of the world, so this is left out of the definition of their value function. This implies that unemployed households derive no benefit from information about η_{s+1} , as it does not affect their choices, and so they choose to process no information. As in the full information model, their Euler equation is unimportant for determining the steady state.

Wealth heterogeneity complicates the solution of the model, as mistakes in saving decisions due to noisy signals are somewhat less costly for wealthier households, which implies they want to pay less for information than poorer households. Both this and the dependence of value functions on prior beliefs mean that households do not return exactly to the same set of prior beliefs each period, as they did in the simple model. Numerically, however, we find that these effects are small,²⁷ so all

²⁵This can be thought of as a timing assumption: households compute the value functions, then use them to solve the rational inattention problem. The solution, as in Section 3.2, comprises an information strategy (what signal to choose) and an action strategy (how to react to each possible signal realisation). The timing assumption is that the household cannot re-evaluate their information or action strategies within the period as a result of the observed signal realisation.

²⁶We assume that households do not infer anything about the true hiring rate from the interest rate. This is consistent with the ignorance prior: households do not know the mapping from other variables to the hiring rate.

²⁷This is a feature of this model. In other settings wealth has been found to have substantial

households choose signals with two discrete realisations, and a high signal realisation induces a similar posterior belief in all households. To make progress, we therefore abstract away from these effects, so as in Section 3.8 households switch between the same two prior beliefs each period. For details of this approximation, and the rest of the solution method, see Appendix E.2.

The steady states of the model are plotted below, for a monthly calibration based on that in Ravn and Sterk (2020) Appendix A1.2 (see our Appendix E.3). As in the static model of Section 3, we consider an information cost ψ such that households optimally limit themselves to a menu with two choices of saving b_{es+1} . The steady state Phillips Curve and Euler Equation (with either full information or rational inattention) pin down the steady state as they both contain just two endogenous variables: inflation and the hiring rate.




Figure 6: Steady state relations under perfect information and rational inattention ($\psi = 0.025$). Calibration in Appendix E.3.

The black (dashed) and red (solid) lines are the Phillips Curve and full information Euler Equation. They match Figure 2 panel 1 in RS, except we have not plotted separate relationships for regions where the zero lower bound on nominal interest rates and the non-negativity constraint on vacancies bind, as we abstract from those constraints. The blue solid line is the steady state Euler Equation with rational inattention. As with the steady state consumption function in Section 3.8, all points along the steady state Euler equation are such that the belief distribution is stationary given aggregate variables (in this case Π as well as η). In addition, the wealth distribution is also stationary given Π and η , and the net asset position effects on information choices (see e.g. Lei, 2019; Macaulay, 2021).

is zero so the asset market clears. This does not mean that wealth or beliefs are static: each household receives signals with idiosyncratic noise, which implies churn in both wealth and beliefs underneath their stable distributions.

The Euler Equation is upward sloping because a higher steady state hiring rate η decreases the desire for precautionary saving.²⁸ 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. The upward slope of the Phillips Curve is a consequence of quadratic price adjustment costs (Ascari and Rossi, 2012).

Under perfect information the interaction of these curves leads to a unique steady state, with $\eta = 0.45$ and $\Pi = 1.001$. With rational inattention there are three steady states, a high employment steady state with $\eta = 0.84$, $\Pi = 1.007$, a middle steady state with $\eta = 0.52$, $\Pi = 1.002$, and a low employment steady state with $\eta = 0.22$, $\Pi = 0.997$.

These steady states share many properties with the simple model in Section 3. The outer two steady states are locally stable with respect to belief mutations while the middle steady state is unstable. Large mutations to beliefs can however shift the economy between the steady states.

Just as in the simple model, underneath each steady state there is heterogeneity and churn in prior beliefs. Households are divided into optimists and pessimists based on the previous realisation of their signal about the hiring rate, and each period some households transition between the two priors.²⁹ This heterogeneity and idiosyncratic variation in beliefs is consistent with the survey evidence presented in the next section.

5 Survey expectations

In the Survey of Consumer Expectations (SCE), 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 expected hiring rate studied in our model and the model of Ravn and Sterk (2020). In this section we show that the survey responses display

²⁸This is because fixed wages imply labour income risk is countercyclical. If labour income risk was procyclical, the Euler equation would slope downwards.

²⁹See Appendix E.4 for details of the prior belief transitions and the effects of belief mutations.

several features which are present in our model. We use the SCE microdata from June 2013 - March 2018.

Figure 7 shows the histogram of responses as deviations from the mean response that month.




Figure 7: Histogram of hiring rate expectations from the Survey of Consumer Expectations, deviations from the average for the month of the interview.

There is a great deal of dispersion. The mean within-month standard deviation of responses is 32%. However, this disagreement could all be due to household heterogeneity, not differences in information. High and low education households, for example, could agree exactly about aggregate labour market conditions, but expect different hiring rates because they are making predictions for different segments of that labour market. To account for as much of this as possible, we run the following regression:

$$\mathbb{E}_{it}\eta_{it,t+3} = \alpha_0 + \alpha_1 X_{it} + \varepsilon_{it} \quad (40)$$

The dependent variable is household i 's expectation of their own hiring rate for the months from t to $t + 3$, and X_{it} contains a wide range of personal characteristics.³⁰ The R^2 for this regression is just 0.084. Adding time fixed effects to pick up any variation due to time-varying aggregate factors only increases this a small amount,

³⁰The controls are age, age², income, income², financial distress, education, gender, race, job tenure, state, home ownership, marital status, 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. All controls except age, age², income, income², and financial distress are categorical or dummy variables. Income is coded as the mid-point of each bin, with the \$200,000+ bin coded as \$250,000.

to 0.095. The vast majority of heterogeneity in labour market expectations does not come from the dimensions of household heterogeneity recorded in the SCE. There are of course other dimensions of household heterogeneity which are not collected in the survey, which could explain more of the heterogeneity, but it is unlikely that this would account for all of the currently unexplained heterogeneity.

Beliefs could, however, be heterogeneous and still co-ordinate on a high (or low) unemployment steady state. If all households agree in some period that the future hiring rate has fallen by 5% relative to their expectations in the previous period, that would give the same self-fulfilling expectations mechanism described in Ravn and Sterk (2020) and Heathcote and Perri (2018), even if the cross-sectional distribution of beliefs features lots of heterogeneity each period. To explore this, we use the panel nature of the SCE to study changes in the unexplained part of expectations, $\Delta\varepsilon_{it}$. If there is strong co-ordination in beliefs, then changes in household-level expectations should be largely explained by changes in average expectations. We therefore decompose individual belief updates into the average belief update that period and an idiosyncratic term, which are orthogonal to each other by construction:

$$\Delta\varepsilon_{it} = \Delta\bar{\varepsilon}_t + u_{it} \tag{41}$$

If households agree on movements in the hiring rate but have different (constant) degrees of optimism or pessimism, the first differencing in this model would remove the individual-specific constants, and the variance of average belief updates ($\Delta\bar{\varepsilon}_t$) would account for most of the variance of individual belief updates ($\Delta\varepsilon_{it}$). That is, the fraction $\sigma_{\Delta\bar{\varepsilon}_t}^2/(\sigma_{\Delta\bar{\varepsilon}_t}^2 + \sigma_{u_{it}}^2)$ would be very close to one. In fact, when we carry out this variance decomposition we find that less than 0.3% of the variation in revisions to household expectations is explained by movements in average expectations. There is therefore a great deal of idiosyncratic variation in *updates* to hiring rate expectations, and only weak co-ordination of beliefs. This is a natural feature of our model, because households form their expectations by observing signals with idiosyncratic noise.

Our model also produces a large amount of persistence in household beliefs through the updating of prior beliefs. This can be seen most clearly in Table 2 above: even in the middle steady state, which features very dispersed expectations, individual households are extremely unlikely to revise their beliefs from one period to the next.³¹ To investigate this persistence in the data, we run the following

³¹This remains true in the HANK model. The equivalent belief transition matrices are in Table 5 in Appendix E.4.

regression:

$$\varepsilon_{it} = \gamma_1 \varepsilon_{it-1} + W_t + e_{it} \quad (42)$$

Here W_t are month fixed effects, which pick up any variation in expectations due to aggregate factors, including average beliefs. The results are displayed in Table 3.

The coefficient γ_1 is large and significant. This implies that hiring rate expectations

Table 3: Regression of the unexplained part of hiring rate expectations on lag of itself and month fixed effects.

| | (1) |
|------------------------|-----------------------|
| | Hiring Rate Residual |
| L.Hiring Rate Residual | 0.720*** (0.00429) |
| Observations | 31604 |
| R^2 | 0.522 |

Robust standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

are indeed persistent at the household level, as predicted by our model.

A final observation from this data is that the variance of expectations is an order of magnitude larger than the variance of the actual (aggregate) quarterly hiring rate, calculated using labour flows in the CPS. Estimating an AR(1) model on the hiring rate, we estimate a long run standard deviation of 8%.³² In contrast, the standard deviation of the unexplained part of hiring rate expectations ε_{it} is 31%. If households knew the true distribution of the hiring rate and received noisy signals (as in a standard rational inattention model a la Sims (2003)), we would expect the variance of beliefs to be *lower* than the variance of the data.³³ 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.³⁴ This supports our assumption that households hold ‘ignorance priors’ about the future hiring rate.

³²This is calculated using labour flows data from 1990-2018. The hiring rate is found to be stationary at the 1% level.

³³This 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 $\mathbb{E}(\eta|s_i)$. The variance of these conditional expectations is given by $V(\mathbb{E}(\eta|s_i)) = V(\eta) - \mathbb{E}(V(\eta|s_i))$. 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.

³⁴This statement relies on our conditioning variables picking up the majority of expectation variation due to heterogeneity in the actual hiring rates experienced by households. Since the variance of expectations is an order of magnitude larger than that of the observed average hiring rate, we are confident that household beliefs are more dispersed than the true distribution of the labour market variables they are trying to estimate.

6 Conclusion

We have proposed a model of a self-fulfilling expectation-driven unemployment trap which is generated by rational inattention. The model features households who face costs of processing information about the future hiring rate, and who do not know the true equilibrium distribution of that rate. The households optimally choose to process noisy signals which imply a highly nonlinear response of aggregate consumption to changes in labour market conditions, which leads to the possibility of multiple steady states. In a HANK model based on Ravn and Sterk (2020) there is a unique steady state if households have full information, but rational inattention generates three steady states: an unemployment trap with a low (but positive) job hiring rate and mild deflation, a middle steady state with higher employment and moderate inflation, and an ‘employment trap’ with very high employment and high inflation. Expectations in the model are consistent with key properties of survey expectations: labour market beliefs are heterogeneous, persistent, and display greater variance than that implied by accurate prior beliefs.

There are numerous directions in which this model framework could be extended. One interesting prospect is how the results change if households pay attention to noisy public signals. In the model as it is, the ignorance prior means that households have no reason to process the information in such public signals beyond what that conveys about the hiring rate. If this information is as costly to process as the private signals we study then households will choose to ignore the public signals. If, however, households have an incentive to learn about the public signal directly (as in Chen et al., 2015) then public signals would play a role, and the way households allocate attention to private and public information could have interesting consequences for the behaviour of the equilibria.

References

- Adam, K., Kuang, P., and Marcet, A. (2012). House Price Booms and the Current Account. *NBER Macroeconomics Annual*, 26(1):77–122.
- Adam, K. and Marcet, A. (2011). Internal rationality, imperfect market knowledge and asset prices. *Journal of Economic Theory*, 146(3):1224–1252.
- Andrade, P., Gaballo, G., Mengus, E., and Mojon, B. (2019). Forward guidance and heterogeneous beliefs. *American Economic Journal: Macroeconomics*, 11(3):1–29.
- Armantier, O., Bruine de Bruin, W., Topa, G., van der Klaauw, W., and Zafar, B.

- (2015). Inflation expectations and behavior: do survey respondents act on their beliefs? *International Economic Review*, 56(2):505–536.
- Ascari, G. and Rossi, L. (2012). Trend Inflation and Firms Price-Setting: Rotemberg Versus Calvo. *The Economic Journal*, 122(563):1115–1141.
- Beaudry, P., Galizia, D., and Portier, F. (2018). Reconciling Hayek’s and Keynes’ Views of Recessions. *The Review of Economic Studies*, 85(1):119–156.
- Carroll, C. D. and Dunn, W. E. (1997). Unemployment Expectations, Jumping (S,s) Triggers, and Household Balance Sheets. *NBER Macroeconomics Annual*, 12:149–217.
- Challe, E., Matheron, J., Ragot, X., and Rubio-Ramirez, J. F. (2017). Precautionary saving and aggregate demand. *Quantitative Economics*, 8(2):435–478.
- Challe, E. and Ragot, X. (2016). Precautionary Saving Over the Business Cycle. *The Economic Journal*, 126(590):135–164.
- Chen, H., Luo, Y., and Pei, G. (2015). Attention misallocation, social welfare and policy implications. *Journal of Economic Dynamics and Control*, 59:37–57.
- Cogley, T. and Sargent, T. J. (2005). The conquest of US inflation: Learning and robustness to model uncertainty. *Review of Economic Dynamics*, 8(2 SPEC. ISS.):528–563.
- Cooper, R. and John, A. (1988). Coordinating Coordination Failures in Keynesian Models. *The Quarterly Journal of Economics*, 103(3):441.
- Diamond, P. A. (1982). Aggregate Demand Management in Search Equilibrium. *Journal of Political Economy*, 90(5):881–894.
- Fox, C. R. and Clemen, R. T. (2005). Subjective probability assessment in decision analysis: Partition dependence and bias toward the ignorance prior. *Management Science*, 51(9):1417–1432.
- Heathcote, J. and Perri, F. (2018). Wealth and volatility. *Review of Economic Studies*, 85(4):2173–2213.
- Jung, J., Kim, J. H. J., Matějka, F., and Sims, C. A. (2019). Discrete Actions in Information-Constrained Decision Problems. *Review of Economic Studies*, 86(6):2643–2667.
- Kandori, M., Mailath, G. J., and Rob, R. (1993). Learning, Mutation, and Long Run Equilibria in Games. *Econometrica*, 61(1):29.

- Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary policy according to HANK. *American Economic Review*, 108(3):697–743.
- Kreps, D. M. (1997). Anticipated Utility and Dynamic Choice. In Jacobs, D. P., Kalai, E., Kamien, M. I., and Schwartz, N. L., editors, *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures*, pages 242–274. Cambridge University Press.
- Lei, X. (2019). Information and Inequality. *Journal of Economic Theory*, 184:104937.
- Luo, Y. (2008). Consumption dynamics under information processing constraints. *Review of Economic Dynamics*, 11(2):366–385.
- Macaulay, A. (2021). The Attention Trap : Rational Inattention , Inequality , and Fiscal Policy. *European Economic Review*, 135:103716.
- Maćkowiak, B., Matějka, F., and Wiederholt, M. (2020). Rational Inattention: A Review. *CEPR Discussion Papers*, no. 15408.
- Maćkowiak, B. and Wiederholt, M. (2009). Optimal Sticky Prices under Rational Inattention. *American Economic Review*, 99(3):769–803.
- Maćkowiak, B. and Wiederholt, M. (2015). Business Cycle Dynamics under Rational Inattention. *Review of Economic Studies*, 82(4):1502–1532.
- Matějka, F. (2016). Rationally Inattentive Seller: Sales and Discrete Pricing. *The Review of Economic Studies*, 83(3):1125–1155.
- Matějka, F. and McKay, A. (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–298.
- Meeks, R. and Monti, F. (2019). Heterogeneous beliefs and the Phillips curve. *Bank of England working papers*.
- Morris, S. and Shin, H. S. (2000). Rethinking Multiple Equilibria in Macroeconomic Modeling. *NBER Macroeconomics Annual*, 15:139–161.
- Ravn, M. O. and Sterk, V. (2020). Macroeconomic Fluctuations with HANK & SAM: an Analytical Approach. *Journal of the European Economic Association*.
- Rotemberg, J. J. (1982). Sticky Prices in the United States. *Journal of Political Economy*, 90(6):1187–1211.
- Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review*, 95(1):25–49.

- Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690.
- Sims, C. A. (2006). Rational inattention: Beyond the linear-quadratic case. *American Economic Review*, 96(2):158–163.
- Steiner, J., Stewart, C., and Matějka, F. (2017). Rational Inattention Dynamics: Inertia and Delay in Decision-Making. *Econometrica*, 85(2):521–553.
- Stevens, L. (2019). Coarse Pricing Policies. *The Review of Economic Studies*, 124(2):420–453.
- Wiederholt, M. (2015). Empirical Properties of Inflation Expectations and the Zero Lower Bound.

A Solving the rational inattention problem

This appendix describes how we solve the rational inattention problem in equations (8)-(11). We begin by discretising the support of η into N equally-spaced points ($N = 30$ for the graphs in Section 3). We then similarly discretise the range of consumption choices into N equally spaced points. The edges of this grid are taken as the minimum and maximum values of c_m^{FI} for $\eta \in [0, 1]$, so the range is the same as that of the optimal consumption that would be chosen if the household had full information about η .

We then define a $N \times N$ matrix $f(\eta, c_m)$, which is the joint distribution of the hiring rate and consumption choices. We can evaluate the expected utility and information costs of any such joint distribution using the formulae in equations (8) and (12). We then use a standard numerical optimisation routine³⁵ to find the matrix $f(\eta, c_m)$ that maximises the expected utility minus the information costs, subject to the constraints that the marginal distribution of η implied by $f(\eta, c_m)$ equals the (discretised) prior $g(\eta)$ (9) and that each element of $f(\eta, c_m)$ is non-negative (10).

The graphs in Section 3 are drawn using $\beta = 0.975$, $R = \frac{1}{\beta}$, $\omega = 0.4$, $w = 1.3$, $\theta = 0.4$, $H(c_m) = -7.67 + 7.75c_m$.

The basic method is the same for the household problem in Section 4, except that the choice variable is saving b_{s+1} , and we do not have an analytic form for the value function. In that case, we incorporate the RI solution into a value function iteration. We guess expected future values of entering a period with a given initial wealth in either labour market state, then use these to form the value functions $\tilde{V}^e(b_s, \eta_{s+1}, b_{es+1}, g(\eta_{s+1}))$ and $\tilde{V}^u(b_s, \eta_{s+1}, g(\eta_{s+1}))$ defined in equations (37)

³⁵fmincon in Matlab.

and (38). We then solve the RI problem, and use the output to update the guesses for the expected values.

As this involves solving the RI problem many times, we speed up the solution by noting that the conditional choice probabilities of any b_{es+1} chosen with positive probability are (Matějka, 2016):

$$f(b_{es+1}^* | \eta_{s+1}, b_s, g(\eta_{s+1})) = \frac{f(b_{es+1}^*) \exp\left(\frac{\tilde{V}^e(b_s, \eta_{s+1}, b_{es+1}^*, g(\eta_{s+1}))}{\psi}\right)}{\sum_{b_{es+1}} f(b_{es+1}) \exp\left(\frac{\tilde{V}^e(b_s, \eta_{s+1}, b_{es+1}, g(\eta_{s+1}))}{\psi}\right)} \quad (43)$$

If we know the levels of b_{es+1} chosen with positive probability and their marginal probabilities $f(b_{es+1})$, we can calculate these conditional probabilities, and use them and the prior $g(\eta_{s+1})$ to find the joint distribution $f(\eta_{s+1}, b_{es+1})$. Rather than solving for each point of the $N \times N$ discretised joint distribution as we do for the simple model, we only need to find the two levels of b_{es+1} that maximise the expected value. We then confirm that the same exercise with three levels implies a lower expected value. This is equivalent to the simpler method employed in Section 3, but gives faster computations.

B Aggregate consumption function in the static model with alternative ψ

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




Figure 8: Aggregate consumption function with $\kappa = 0$ (green), $\kappa = 0.5$ (blue), $\kappa = 1$ (black) and in the unconstrained case (red). Other parameters in Appendix A.

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 flatter regions in the aggregate response curve.

C Static model: firms

Here we microfound the linear equilibrium relationship between aggregate consumption and the hiring rate in Section 3.1.

In the morning, all households are employed at wage 1 and the firm sells \bar{c} units of the consumption good. Assume that any unsold output in the morning is wasted, and that the firm receives some subsidy for operating S . Firm profits are therefore equal to $\bar{c} - 1 + S$.

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

$$h = mv^{1-\alpha}u^\alpha \quad (44)$$

Now note that the number of job seekers at the start of the afternoon is equal to the number of separations at the end of the morning, ω , since all households were employed in the morning. 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:³⁶

$$C(\eta) = km^{\frac{-1}{1-\alpha}} \omega \eta^{\frac{1}{1-\alpha}} \quad (45)$$

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

$$D(\eta) = F - w - C(\eta) \quad (46)$$

Where F is the value of production from that worker, equal to afternoon aggregate consumption per worker plus the value of inventory at the end of the afternoon. Assume that this is sufficiently high that $D(\eta) > 0$ for all $\eta \in (0, 1)$. That is, in the afternoon, 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 morning profits to pay for vacancies.

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

$$\bar{c} - 1 + S = km^{\frac{-1}{1-\alpha}} \omega \eta^{\frac{1}{1-\alpha}} \quad (47)$$

In the graphs in Section 3.1 we 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} and η :

$$\bar{c} = 1 - S + \frac{k\omega}{m} \eta \quad (48)$$

This is the linear relationship used in Section 3.1. The parameters used there, in addition to those given in Appendix E.3, are $m = 0.25$, $k = 0.062w$, $S = 0.01$.

D Dynamic model proofs

D.1 Proof that $f(\eta_2|c_{j,2}) = f(\eta_1|c_{j,1})$

Decomposing the joint distribution $f(c, \eta)$ into $f(c)f(\eta|c)$, substituting constraint (11) into the objective function, and dropping terms independent of policy, we can

³⁶We are assuming that firms are large, so the proportion of vacancies filled equals the probability of filling a vacancy.

write the household problem as:

$$\max_{\hat{f}(c), \hat{f}(\eta|c)} \int_c \int_\eta V(c, \eta) \hat{f}(c) \hat{f}(\eta|c) d\eta dc + \psi \int_c \int_\eta \hat{f}(c) \hat{f}(\eta|c) \log \hat{f}(\eta|c) d\eta dc \quad (49)$$

subject to

$$\int_\eta \hat{f}(\eta|c) d\eta = 1 \quad \forall c \quad (50)$$

$$\int_c \hat{f}(c) \hat{f}(\eta|c) dc = g(\eta) \quad \forall \eta \quad (51)$$

$$\hat{f}(\eta|c) \geq 0 \quad \forall \eta, c \quad (52)$$

$$\hat{f}(c) \geq 0 \quad \forall c \quad (53)$$

The only place in which the prior belief enters the problem is constraint (51). We know that if $g(\eta) = 1$ then the solution has $f(c) > 0$ for two levels of consumption only, which we label c_L and c_H . This proof will proceed to show that the same two levels of consumption, and the same associated posterior beliefs, are also a solution to the household problem when $g(\eta) = \rho f(\eta|c_j) + (1 - \rho)$, with $c_j \in \{c_L, c_H\}$.

By definition, we can write:

$$f(\eta|c_j) = \int_c \delta(c - c_j) f(\eta|c) dc \quad (54)$$

Where $\delta(\cdot)$ is the Dirac delta function.

Supposing that the period 2 posteriors are indeed the same as those from period 1, we can further write that $f(\eta|c) = \hat{f}(\eta|c)$. Substituting (54) into (51), using $f(\eta|c) = \hat{f}(\eta|c)$, and rearranging we obtain:

$$\int_c \left(\frac{\hat{f}(c) - \rho \delta(c - c_j)}{1 - \rho} \right) \hat{f}(\eta|c) dc = 1 \quad (55)$$

That is, the constraint is the same as in the case with a uniform prior, with a transformation of the marginal densities to $f^*(c)$, where:

$$f^*(c_j) = \frac{\hat{f}(c_j) - \rho}{1 - \rho} \quad (56)$$

Since the household objective function is multiplicative in $\hat{f}(c)$, we can therefore write the problem with updated priors exactly as in the uniform case, but with $f^*(c)$ replacing $\hat{f}(c)$. Since c_L and c_H are the optimal menu in the uniform case, they also solve the household problem with the updated prior. The posteriors will

also be the same. This verifies our guess that:

$$f(\eta_2|c_{j,2}) = f(\eta_1|c_{j,1}) \quad (57)$$

To solve for the marginal probabilities with the updated prior, we use the fact that the marginal probabilities of c_L and c_H in the uniform case are equal to 0.5. Substituting this into the definition of $f^*(c_j)$ and rearranging gives:

$$\hat{f}(c_j) = 0.5(1 + \rho), \quad \hat{f}(c_{-j}) = 0.5(1 - \rho) \quad (58)$$

D.2 Steady state local stability

The proportion of households who hold pessimistic priors evolves according to:

$$\mathbf{P}_{t+1} = \phi(\mathbf{P}_t, \eta_t(\mathbf{P}_t)) = \mathbf{P}_t f(c_{L,t}|\eta_t, c_{L,t-1}) + (1 - \mathbf{P}_t) f(c_{L,t}|\eta_t, c_{H,t-1}) \quad (59)$$

Therefore:

$$\frac{d\mathbf{P}_{t+1}}{d\mathbf{P}_t} = f(c_{L,t}|\eta_t, c_{L,t-1}) - f(c_{L,t}|\eta_t, c_{H,t-1}) + \frac{d\eta_t}{d\mathbf{P}_t} \frac{\partial \mathbf{P}_{t+1}}{\partial \eta_t} \quad (60)$$

First, suppose that there is no response of the hiring rate to mutations (set $\frac{d\eta_t}{d\mathbf{P}_t} = 0$). Then, $\frac{d\mathbf{P}_{t+1}}{d\mathbf{P}_t} = \phi_{\mathbf{P}}(\mathbf{P}_t, \eta_t(\mathbf{P}_t)) = f(c_{L,t}|\eta_t, c_{L,t-1}) - f(c_{L,t}|\eta_t, c_{H,t-1})$. The probability of a household choosing c_L in period t is always higher if they also chose it in period $t - 1$ (see Section 3.8). Therefore $\phi_{\mathbf{P}}(\mathbf{P}_t, \eta_t(\mathbf{P}_t)) \in (0, 1)$, so any belief distribution is locally stable.

If we allow the hiring rate to vary, however, it is possible for a steady state to be unstable. From here we simplify notation by denoting $f_{jk} = f(c_{j,t}|\eta_t, c_{k,t-1})$, and $f'_{jk} = \frac{\partial f_{jk}}{\partial \eta_t}$.

First, note that:

$$\frac{\partial \mathbf{P}_{t+1}}{\partial \eta_t} = \mathbf{P}_t f'_{ll} + (1 - \mathbf{P}_t) f'_{lh} \quad (61)$$

The hiring rate always satisfies the firm condition $\eta_t = H(\bar{c}_t)$, and aggregate consumption satisfies (19), which simplifies to:

$$\bar{c}_t = c_H + (c_L - c_H)\mathbf{P}_{t+1} \quad (62)$$

We therefore have:

$$\frac{d\eta_t}{d\mathbf{P}_t} = H'(\bar{c}_t) \frac{d\bar{c}_t}{d\mathbf{P}_t} = H'(\bar{c}_t)(c_L - c_H) \frac{d\mathbf{P}_{t+1}}{d\mathbf{P}_t} \quad (63)$$

Substituting this and (61) in to (60) and rearranging gives:

$$\frac{d\mathbf{P}_{t+1}}{d\mathbf{P}_t} = \frac{f_u - f_{lh}}{1 - H'(\bar{c}_t)(c_L - c_H)(\mathbf{P}_t f'_{lu} + (1 - \mathbf{P}_t)f'_{lh})} \quad (64)$$

Assuming that $H'(\bar{c}_t)(c_L - c_H)(\mathbf{P}_t f'_{lu} + (1 - \mathbf{P}_t)f'_{lh}) < 1$,³⁷ we have instability if:

$$H'(\bar{c}_t)(c_L - c_H)(\mathbf{P}_t f'_{lu} + (1 - \mathbf{P}_t)f'_{lh}) > 1 - f_u + f_{lh} \quad (65)$$

The slope of the firm condition in Figure 5 is $(H'(\bar{c}_t))^{-1}$. The slope of the steady state aggregate consumption function is:

$$\frac{dc^{ss}(\eta)}{d\eta} = \frac{d}{d\eta}(c_H + (c_L - c_H)\mathcal{P}(\eta)) = (c_L - c_H)\frac{d\mathcal{P}(\eta)}{d\eta} \quad (66)$$

Where:

$$\frac{d\mathcal{P}(\eta)}{d\eta} = \frac{d}{d\eta}\left(\frac{f_{lh}}{1 - f_u + f_{lh}}\right) = \frac{f'_{lu}\mathcal{P}(\eta) + f'_{lh}(1 - \mathcal{P}(\eta))}{1 - f_u + f_{lh}} \quad (67)$$

Evaluating condition 65 at steady state (so $\mathbf{P}_t = \mathcal{P}(\eta_t)$) gives that a steady state is unstable if the steady state aggregate consumption function is steeper than the firm condition:

$$\frac{dc^{ss}(\eta_t)}{d\eta_t} > (H'(\bar{c}_t))^{-1} \quad (68)$$

This means that the unemployment trap and employment trap steady states are locally stable, because the steady state aggregate consumption function is shallower than the firm condition (see Figure 5). The middle steady state, however, is unstable, because the steady state aggregate consumption function is steeper than the firm condition.³⁸

To see the intuition for this, consider a mutation that reduces aggregate consumption by a given amount. The slope of the steady state aggregate consumption

³⁷This is the necessary condition for $\frac{d\eta_t}{d\mathbf{P}_t} < 0$. It can be shown that this always holds at the outer steady states, and indeed it is also true at the middle steady state for most calibrations (including the one used to draw the figures in Section 3). Following the same steps as this proof, it can be shown that if $\frac{d\eta_t}{d\mathbf{P}_t} > 0$, a steady state is unstable if $\frac{dc^{ss}(\eta_t)}{d\eta}(1 - f_u + f_{lh}) < (H'(\bar{c}_t))^{-1}(1 + f_u - f_{lh})$. Since $\frac{d\eta_t}{d\mathbf{P}_t} > 0$ implies $\frac{dc^{ss}(\eta_t)}{d\eta}(1 - f_u + f_{lh}) > (H'(\bar{c}_t))^{-1}$, this type of steady state is unstable when $f_u - f_{lh}$ is large, which occurs when η is close to 0.5, as is the case for the middle steady state in a wide range of calibrations.

³⁸Note that this does not necessarily mean that the middle steady state is unstable with respect to all shocks. We study these kinds of mutation shocks because they leave the household information and choice problems, the set of prior beliefs, and the firm side of the model identical, which makes them tractable. This will not generally be true of other shocks. Movements in the firm equilibrium condition, or other aspects of the household problem than simply the priors, will cause shifts in either or both of the two curves defining equilibrium, and it is not clear whether the middle steady state would remain unstable with respect to those kinds of disturbances.

function gives how far the hiring rate would have to fall to sustain this new lower \bar{c} with beliefs at their stationary distribution: if $\bar{c}^{ss}(\eta)$ is steeper, η needs to fall by less for any given \bar{c} fall. The firm condition then gives how far the hiring rate actually falls in equilibrium. If it is shallower than $\bar{c}^{ss}(\eta)$, the hiring rate falls by more than would be necessary to sustain the mutated beliefs as a stationary distribution. This means that the proportion of pessimists will be even higher in period $t + 1$ than in period t , and the original steady state will be unstable. In contrast, if the firm condition is shallower than $\bar{c}^{ss}(\eta)$, then η_t falls by less than would be required to sustain \bar{c}_t as a steady state. The hiring rate is therefore insufficiently responsive to the mutation to create instability.

D.3 Large mutations

Suppose we begin at a steady state with a hiring rate of η_{ss} . If a mutation decreases the proportion of pessimists to $\mathbf{P}_t < \mathcal{P}(\eta_{ss})$, we show that in the absence of other shocks:

$$\mathbf{P}_{t+1} < \mathbf{P}_t \quad \text{if } \bar{c}^{ss}(\eta_t) > H^{-1}(\eta_t) \quad (69)$$

That is, the proportion of pessimists falls even further in the following period if the steady state aggregate consumption function is above the firm condition at the equilibrium hiring rate in the period of the mutation. If this is true, then the steady state is unstable with respect to this large mutation: this is simply an analogue of the result in Appendix D.2 for large shocks. Similarly, if a mutation increases \mathbf{P}_t , $\mathbf{P}_{t+1} > \mathbf{P}_t$ if $\bar{c}^{ss}(\eta_t) < H^{-1}(\eta_t)$.

First, notice that the mapping in equation (59) implies:

$$\mathbf{P}_{t+1} = \mathbf{P}_t f(c_{L,t}|\eta_t, c_{L,t-1}) + (1 - \mathbf{P}_t) f(c_{L,t}|\eta_t, c_{H,t-1}) = \mathcal{P}(\eta_t) \quad (70)$$

That is, absent further mutations in period $t + 1$, the proportion of pessimists in period $t + 1$ is equal to the proportion in the stationary belief distribution at η_t , the equilibrium hiring rate implied by the mutated proportion of pessimists \mathbf{P}_t .

We next rearrange (19) to express \mathbf{P}_t as a function of \bar{c}_t , c_L , c_H , and conditional choice probabilities. Similarly, rearrange (21) to express $\mathcal{P}(\eta_t)$ in terms of $\bar{c}^{ss}(\eta_t)$, c_L , and c_H . Using these expressions, we can write the statement $\mathbf{P}_{t+1} < \mathbf{P}_t$ as:

$$\begin{aligned} \bar{c}_t - c_H(1 - f(c_{L,t}|\eta_t, c_{L,t-1})) - c_L f(c_{L,t}|\eta_t, c_{H,t-1}) \\ < \bar{c}^{ss}(\eta_t)(f(c_{L,t}|\eta_t, c_{L,t-1}) - f(c_{L,t}|\eta_t, c_{H,t-1})) \end{aligned} \quad (71)$$

Using the definitions of $\bar{c}^{ss}(\eta_t)$ (21) and $\mathcal{P}(\eta_t)$ (20) this simplifies to:

$$\bar{c}_t < \bar{c}^{ss}(\eta_t) \quad (72)$$

Similarly, we have that:

$$\mathbf{P}_{t+1} > \mathbf{P}_t \quad \text{if } \bar{c}_t > \bar{c}^{ss}(\eta_t) \quad (73)$$

In all of these equations, the firm condition implies that $\bar{c}_t = H^{-1}(\eta_t)$.

If a mutation that reduces the proportion of pessimists arrives when the economy is at the unemployment trap steady state in Figure 5, then condition (72) fails so long as the aggregate consumption function is below the firm condition at the resulting hiring rate. If the mutation is sufficiently large that η_t is greater than the hiring rate associated with the middle steady state, then (72) will be satisfied and the proportion of pessimists will continue to fall with no further shocks, eventually leading to the employment trap steady state. Denoting the hiring rate at the middle steady state as η_m , the economy will move from the unemployment trap to the employment trap for any shock that implies $\eta_t > \eta_m$. Similarly, a pessimistic mutation at the employment trap will cause a shift to the unemployment trap if the mutation is large enough to push $\eta_t < \eta_m$.

Our assumption of inertia plays a significant role in determining the size of mutation required to achieve this. Although escaping the unemployment trap only requires the hiring rate to rise an arbitrarily small amount above η_m , this does not mean that $\mathbf{P}_t < \mathcal{P}(\eta_m)$ is always sufficient to cause the transition, and similarly a mutation to a belief distribution with more pessimists than seen at the middle steady state will not necessarily cause a transition from the employment trap to the unemployment trap.

To see why, consider a large mutation arriving when the economy is at the employment trap steady state, such that the proportion of pessimists increases to exactly $\mathcal{P}(\eta_m)$. Since this is the belief distribution associated with the middle steady state, $\eta_t = \eta_m$ is one possible equilibrium in period t . With the calibration used to draw Figure 5, this is at $\eta_t = 0.49$. However, that is not the only hiring rate at which aggregate consumption is consistent with the firm condition. There are two other possible equilibria, with $\eta_t = 0.69$ and $\eta_t = 0.87$. The inertia assumption means we select $\eta_t = 0.87$, as this is closest to the employment trap steady state hiring rate of the previous period. Most households therefore receive a signal that η is high, and so in the next period the proportion of pessimists is much lower than in period t . Condition (73) is not satisfied, and the economy returns to the employment trap. To transition to the unemployment trap the mutation must therefore be sufficiently

large that there is no such high employment equilibrium in period t , which for these parameters occurs when \mathbf{P}_t rises to 0.88.³⁹

The belief distribution at which a given equilibrium ceases to exist is a function of the parameters of the model. It is therefore possible in principle to choose parameters such that the middle steady state is precisely at the boundary, so for any mutation in a given direction the associated equilibrium ceases to exist. In this case the hiring rate is not continuous in the proportion of pessimists and the results of Appendix D.2 are not applicable. The results presented here on large mutations apply in these knife-edge cases.

E HANK model details

E.1 Firm problem

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

$$\max_{P_{js}, v_{js}} \mathbb{E}_t \sum_{s=t}^{\infty} \Lambda_{j,t,t+s} \left[\frac{P_{js}}{P_s} y_{js} - w_s n_{js} - k v_{js} - \frac{\phi}{2} \left(\frac{P_{js} - P_{j,s-1}}{P_{j,s-1}} \right)^2 \right] \quad (74)$$

subject to

$$y_{js} = \left(\frac{P_{js}}{P_s} \right)^{-\gamma} y_s \quad (75)$$

$$y_{js} = e^{A_s} n_{js} \quad (76)$$

$$n_{js} = (1 - \omega) n_{j,s-1} + v_{js} q_s \quad (77)$$

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

Solving this profit maximisation we have:

$$1 - \gamma + \gamma m c_s = \phi (\Pi_s - 1) \Pi_s - \phi \mathbb{E}_s \Lambda_{s,s+1} \left[\frac{y_{s+1}}{y_s} (\Pi_{s+1} - 1) \Pi_{s+1} \right] \quad (78)$$

³⁹With different parameters it can be the case that $\mathbf{P}_t > \mathcal{P}(\eta_m)$ is sufficient to leave the employment trap. Indeed with the calibration used for Figure 5, $\mathbf{P}_t < \mathcal{P}(\eta_m)$ is sufficient to escape the unemployment trap.

where

$$mc_s = \frac{w_s}{e^{A_s}} + \frac{k}{q_s} - (1 - \omega) \mathbb{E}_s \Lambda_{s,s+1} \frac{k}{q_{s+1}} \quad (79)$$

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

The labour market matching function is:

$$M(u_s, v_s) = u_s^\alpha v_s^{1-\alpha} \quad (80)$$

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

$$q_s = \eta_s^{\frac{-\alpha}{1-\alpha}} \quad (81)$$

Firms are owned by the asset-rich, who are risk neutral, so the stochastic discount factor $\Lambda_{s,s+1} = \beta$ for all periods s . The technology shock A_s equals 0 in steady state. The steady state Phillips Curve is therefore:

$$\phi(1 - \beta)(\Pi - 1)\Pi = 1 - \gamma + \gamma(w + k\eta^{\frac{\alpha}{1-\alpha}}(1 - \beta(1 - \omega))) \quad (82)$$

E.2 Solution method

Here we detail the method used to find the steady state Euler equation under rational inattention in Figure 6. We begin by normalising current prices P_s to 1, and fixing inflation at some level Π . We then solve the rational inattention problem of an employed household with a uniform prior belief and zero net wealth ($b_s = 0$) (see Section 4.3 for the details of this problem, and Appendix A for the solution method). As in the simple model of Section 3, households choose signals with two discrete realisations. As in Section 3.8, we take the two posteriors from the solution and update them to form next-period priors using equation (16). As in the diagrams in Section 3.8, we use a persistence of $\rho = 0.9$. We then solve the household problem again for each of these prior beliefs, again for a household with $b_s = 0$.

In the simple model in Section 3.8, the posteriors formed with the updated prior beliefs are the same as in the initial period with the uniform prior for all households. While priors retain their logistic shape (the results from Matějka (2016) still apply), this is no longer exact in the HANK model, for two reasons. First, the updating of prior beliefs affects the continuation value of savings. Households therefore face a different value function in the initial period than in subsequent periods when they have updated their priors. Second, $b_s = 0$ is not the only possible level of initial wealth. A household receiving a signal that η_{s+1} is low in the initial period will save more than a household receiving a signal that η_{s+1} is high. Since wealthy

households face a somewhat lower cost of mistakes than poorer households, if we left the model as it is wealthy households would process a little less information than poorer households. While all households with the wealth levels seen in steady state in our calibration still choose a signal with two possible realisations, the extent to which beliefs are updated from their priors, and so the shape of the prior beliefs in the next period, would therefore depend on the household's wealth history as well as their previous signals.

Numerically, however, these differences are small. As an example, consider the results when $\Pi = 1.0025$. If two households with $b_s = 0$ receive a signal that η_{s+1} is low, but they have the two different priors associated with each of the two signal realisations from the uniform initial prior, the household with the high prior has $\mathbb{E}_s \eta_{s+1} = 0.259$ and the household with the low prior has $\mathbb{E}_s \eta_{s+1} = 0.263$. Furthermore, we find that numerically the signals chosen by the two households imply the same menu of savings choices relative to the decision rule of a household who knows η_{s+1} for certain.⁴⁰ Similarly, if a third household also has the prior associated with a low signal from the uniform prior, but they have $b_s = 0.03$,⁴¹ then their low signal implies $\mathbb{E}_s \eta_{s+1} = 0.280$.

We therefore proceed by making two approximations to abstract from these mechanisms. First, we abstract from the effects of wealth on information processing by assuming that wealth can only enter household decisions through the *action strategy*, not through the *information strategy*. That is, all households choose signals as if they had zero assets. When they observe a realisation of the signal, however, they respond to their newly updated hiring rate expectations taking account of their wealth. Households that start a period with more wealth will choose to save more than a poorer household with the same beliefs.

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 belief about η , which it reports

⁴⁰As described in Appendix A, we set the range of values of household choices (in this case b_{es+1}) for the rational inattention problem as the range of values that would be chosen by a fully-informed household facing the same value function. The different priors of the two households imply different continuation values, and so the equivalent fully-informed households choose different decision rules (they are only fully-informed about the variable the inattentive household is learning about, so not about values of the hiring rate in period $s + 2$ and onwards). Although the two inattentive households choose a different menu of saving choices, we find that the menus are the same relative to the ranges set by the decision rules of these corresponding fully-informed households.

⁴¹This would put them in the top 27.5% of the wealth distribution at the middle steady state, where belief and wealth dispersion are very high relative to the other steady states. A wealth of 0.03 puts a household in the top 0.1% of the wealth distribution at the unemployment trap and employment trap steady states.

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).⁴²

Wealth and information are not 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 restriction removes is that wealth, in turn, affects the optimal amount of information processed, and so affects beliefs in future periods. The numerical results above suggest that this link is weak, implying that the signals themselves do not differ very much by wealth, even though actions in response to those signals do.

Our second approximation is that if a household receives a signal that η_{s+1} is low, we set their prior for the following period to the low prior from period 2, after one round of updating from the uniform prior. Similarly, if a household receives a high signal, we set their subsequent prior equal to the high prior formed after one round of belief updating. These approximations imply that households switch between a set of two prior beliefs, as they do in the simple model. As in Section 3.8, we refer to these priors as optimistic and pessimistic.

The resulting joint distributions $f(\eta_{s+1}, b_{s+1})$ give, for each prior belief and initial wealth b_s , the transition probabilities between beliefs and wealth levels at each level of η_{s+1} . We discretise the range of possible wealth levels into a grid of 501 points spaced equally on $[-0.1, 0.1]$. This covers the vast majority of wealth levels seen in steady state, with less than 0.15% of households at the boundaries at all points along the steady state Euler equation. We find the stationary joint distribution of beliefs and wealth for each level of the hiring rate using the transition probabilities from the joint distributions, adjusted for the fact that a fraction ω of workers become unemployed each period, to be replaced by households newly hired from unemployment, who have the same belief distribution as the rest of the employed population, but who all have $b_s = 0$.⁴³ Summing over the wealth distribution gives the net asset position of employed households. This declines monotonically as the hiring rate rises, because the probability of receiving high signals rises, which en-

⁴²Computationally, we implement this by assuming that the joint distribution $f(\eta_{s+1}, b_{s+1})$ for a household with wealth $b_s = b$ is equal to that found for the household with $b_s = 0$, with the b_{s+1} axis rescaled according to the range of b_{s+1} choices made by a fully-informed household with the same priors and wealth $b_s = b$.

⁴³The prior belief distribution of newly hired households is identical to that of the population as a whole in steady state because no information processing takes place while the household is unemployed, so their priors remain as they were when they left employment.

courages households to save less. There is therefore at most one hiring rate at which the net assets of employed households are zero, as required for the asset market to clear. If there is no such point, the chosen level of inflation is never consistent with a steady state of the Euler equation. If there is, we select the hiring rate consistent with asset market clearing. The resulting pair (η, Π) is a co-ordinate on the steady state Euler equation. Repeating this process for a range of values for Π allows us to trace out the steady state Euler equation plotted in Figure 6.

E.3 Calibration

Table 4 lists the parameter values used to draw Figure 6 in Section 4 and Figure 9 in Appendix E.4. The calibration is monthly, and is based on Ravn and Sterk (2020) Appendix A1.2. We set the annualised inflation target to 2%, the nominal interest rate target such that the annualised real interest rate is 2% when inflation is at target, and we choose the price stickiness parameter ϕ such that the linearised Phillips Curve in a full-information, complete-insurance version of the model would have the same slope as the Phillips Curve in an equivalent model with Calvo price setting and an average price duration of 6 months. The parameter ψ , the marginal cost of information, is set to ensure that each household always chooses an optimal menu of savings choices with two discrete savings levels. We choose the wage so that the full information steady state has a monthly hiring rate of 0.45, as in Shimer (2005). Finally, the labour disutility parameter ζ is set such that the household is indifferent between employment and unemployment under full information. The wage and home production value are therefore consistent with Nash bargaining when the worker has no bargaining power (see Ravn and Sterk (2020) Appendix A1 for details).

Table 4: HANK model calibration

| Parameter | Name | Value |
|-------------|--------------------------------------|-----------------------|
| γ | Elasticity of substitution | 6 |
| ϕ | Price adjustment cost | 146.3 |
| k | Vacancy posting cost | $0.1w$ |
| α | Elasticity of matching function | 0.5 |
| ω | Separation rate | 0.02 |
| w | Real wage | 0.83252 |
| θ | Home production value | $0.8w$ |
| β | Discount factor | 0.995 |
| μ | Coefficient of risk aversion | 2 |
| ζ | Disutility of labour | $(4w)^{-1}$ |
| ψ | Information processing cost | 0.025 |
| $\bar{\Pi}$ | Monthly inflation target | $1.02^{\frac{1}{12}}$ |
| \bar{R} | Monthly nominal interest rate target | $1.02^{\frac{1}{24}}$ |
| d_{Π} | Taylor Rule parameter | 1.5 |

E.4 Transitions

At the individual level there are transitions between different prior beliefs every period, even when the aggregate economy is in steady state. As in Section 3.8, in steady state households have either optimistic or pessimistic prior beliefs, and while the proportions of these are constant at each of the steady states individual households churn between the two priors. Table 5 gives the transition probabilities between the two prior beliefs at each steady state:

Table 5: Prior belief transition matrices

| (a) Unemployment trap | | | (b) Middle steady state | | | (c) Employment trap | | |
|-----------------------|-----------|-----------|-------------------------|-----------|-----------|---------------------|-----------|-----------|
| | P_{s+1} | O_{s+1} | | P_{s+1} | O_{s+1} | | P_{s+1} | O_{s+1} |
| P_s | 0.9993 | 0.0007 | P_s | 0.9749 | 0.0251 | P_s | 0.4400 | 0.5600 |
| O_s | 0.9149 | 0.0851 | O_s | 0.0531 | 0.9469 | O_s | 0.0002 | 0.9998 |


At the aggregate level, mutations to the distribution of prior beliefs, as studied for the simple model in Section 3.9, can shift the economy between steady states. To see this, we run a simulation of the model with mutations. At the start of each simulation period, we draw two ‘mutation’ shocks. We first take a fraction m_p of households with pessimistic priors and mutate their beliefs so they become optimistic. The mutating households are drawn at random, so over the large popu-

lation of households we take the same proportion of the pessimistic households from each point in the wealth distribution. Simultaneously, we take a proportion m_o of households with optimistic priors and make them pessimistic.


We draw m_p, m_o from an i.i.d. truncated generalised Pareto distribution (location=0, scale=0.1, shape=50, truncated to be in the range $[0, 1]$). We use a fat-tailed distribution to ensure that, on occasion, the system does transition from one steady state to another - but that most of the time shocks will be small enough for the economy to remain in the neighbourhood of the steady state.

Over many periods of random mutation shocks, the hiring rate displays a bimodal distribution, spending most of the time close to the high employment steady state, but occasionally transitioning and spending a substantial number of periods around the unemployment trap before a sufficiently large mutation shifts the economy back to the high steady state. This is shown in Figures 9a and 9b .

Figure 9: Results from hiring rate simulations



(a) Histogram of η realisations over a 10000 period simulation.



(b) η over simulation periods 1-100. Red dotted lines show steady states.