

Cyclical Attention to Saving*

Alistair Macaulay[†]

August 25, 2020

Abstract

This paper explores the business cycle implications of limited household attention to choosing between different savings products. In a model with heterogeneous banks, savers pay more attention to their bank choice when the marginal utility of income is high. This implies that attention rises in contractions. I find evidence for such countercyclical attention using a novel combination of data on retail savings markets in the UK. In the data, banks offer heterogeneous interest rates on very similar products, and savers more reliably choose products closer to the top of the rate distribution during contractions. Countercyclical attention amplifies shocks to consumption: after a contractionary shock, attention rises, so savers experience higher interest rates, which causes a further fall in consumption. In a quantitative New Keynesian model, this amplification is estimated to be large. Countercyclical attention increases the variance of consumption by 17%, and amplifies some key shocks by more than 25%. Policies that reduce the costs of comparing between financial products have substantial stabilization effects.

JEL codes: D83, E32, E71

*I thank Klaus Adam, Julio A. Blanco, Vasco Carvalho, Wei Cui, Martin Ellison, Jesus Fernandez-Villaverde, Federico Kochen, Tobias König, John Leahy, Sang Lee, Filip Matějka, Michael McMahon, Christopher J. Palmer, Luigi Pistaferri, Ricardo Reis, Matthew Shapiro, Boromeus Wanengkirtyo, Mirko Wiederholt, and seminar and conference participants at the 3rd Annual NuCamp Conference, Applied Young Economist Webinar series, EEA Annual Congress 2020, Econometric Society World Congress 2020, Royal Economic Society Symposium of Junior Researchers 2019, University of Birmingham, University of Michigan, University of Oxford, and the Young Economists Symposium 2020 for helpful comments and suggestions. I am particularly grateful to Richard Harrison for kindly sharing the code for Harrison and Oomen (2010), and to Moneyfacts for permission to use their data.

[†]Department of Economics & Nuffield College, University of Oxford. Email: alis-tair.macaulay@economics.ox.ac.uk

1 Introduction

In the majority of dynamic macroeconomic models the interest rate is crucial in determining how shocks propagate through the economy, in part because it regulates the consumption of intertemporally maximising households. The interest rate is usually taken as given by households in these models, but regulators have noted that in reality savers face a range of rate-bearing products, and that they could increase the interest rate they earn on their savings by ‘shopping around’ for the best product (FCA, 2015).

In this paper I ask if the extent of shopping around, or *attention* to product choice, varies systematically with the business cycle. I find in both theory and data that attention is countercyclical. This substantially amplifies shocks in an estimated business cycle model, because of the effects of attention on the interest rate that households experience.

I first develop a simple model to explore the interaction of rationally inattentive savers¹ and deposit-taking banks. Profit-maximising banks face heterogeneous costs, and in the face of incomplete attention from savers they offer heterogeneous interest rates. If a household pays more attention, they increase their probability of choosing a bank offering a high interest rate, and so they increase the average interest rate they face. The key drivers of attention in this environment are the marginal utility of future income and the extent of interest rate dispersion.

The marginal utility of income drives the countercyclical behaviour of attention, which in turn implies that variable attention amplifies business cycle fluctuations. Consider, for example, a shock that causes consumption to fall. The marginal utility of income rises, and so households pay more attention to their choice of savings product: intuitively, it becomes more important to extract every possible dollar of interest income out of their savings, and so they pay more attention in order to achieve that. That means they face higher interest rates relative to the distribution of offered rates. In addition, with all savers paying more attention the deposit market is more competitive, causing banks to offer higher interest rates. Through two channels, the household therefore faces higher interest rates than they would have done if attention had stayed constant, and higher interest rates cause consumption to fall even further through a standard consumption Euler equation. Countercyclical attention therefore amplifies the consumption fall.

I find evidence of countercyclical attention to savings using a novel combination of data on savings markets in the UK. Detailed product-level data reveals substantial dispersion in the interest rates banks offer on a set of extremely similar products at any point in

¹I follow the discrete choice rational inattention framework of Matějka and McKay (2015). I show in appendix A.1 that a similar model with endogenous search effort without congestion externalities, as in McKay (2013), gives the same qualitative implications as in the main body of the paper. The extension of persistent interest rates (see appendix C.1) would however be intractable in this alternative setup.

time. Linking this with data on the average interest rates achieved by savers opening new products in this set, I show that savers on average choose products higher up in the interest rate distribution in contractions, as predicted by the model. For this analysis I focus on fixed interest rate products, as their simplicity gives me the best chance of ruling out that rate dispersion and saver decisions are being driven by unobservable product differentiation. This should be viewed as a useful laboratory in which to study household behaviour: none of the mechanisms I explore are specific to this market, or to the UK².

The existence of interest rate dispersion is an important prerequisite for attention to affect the interest rate households face. I obtain panel data on the savings products available in the UK by digitising monthly editions of Moneyfacts, a magazine for UK financial advisers. There is substantial dispersion in offered interest rates even among products which are identical across the wide range of product features reported. Considering institutional details of the UK savings market, I argue that unobserved product heterogeneity is unlikely to explain the majority of this dispersion. Instead, I argue that much of this interest rate dispersion persists in equilibrium because of an information friction: it is costly for households to acquire information about the set of products on offer. The existence of the Moneyfacts data is itself a justification for the information cost interpretation. Financial advisers (and indeed the Bank of England and several other regulators) would not need to pay for such a magazine if the information was easy to obtain elsewhere³.

The model predicts that savers should experience higher interest rates relative to this distribution of offered rates in contractions, as they increase attention. This is precisely what I find in the data. Data from the Bank of England gives the average interest rate achieved on new accounts opened each month for specific sets of savings products with particular characteristics. Identifying the set of products with those characteristics in the Moneyfacts data, I find that the position of the rate households achieve within the distribution of offers is indeed countercyclical. When the unemployment rate is high, and the level of average interest rates in the market is low⁴, households on average choose products that are higher up within the distribution of interest rates.

To quantify the importance of countercyclical attention for shock transmission, I build a medium-scale small open economy New Keynesian model of the UK based on that of

²The mechanism through which variable attention to savings products affects the business cycle does not necessarily apply in the same way to loans. In addition, I only consider the effects of countercyclical attention to savings on consumers, leaving aside the question of what this means for the allocation of credit. I discuss these additional channels further in appendix A.2.

³Staff at the regulator found that shopping around decisions were indeed driven by an analysis of the costs and benefits, including time spent shopping and likely interest rate gains (Cook et al., 2002).

⁴In the estimated model I allow for distinct channels through which each of these could affect attention. I find that for most contractionary supply shocks - after which unemployment and interest rates both rise - attention increases. The marginal utility of income is the dominant driver of attention.

Harrison and Oomen (2010), featuring many of the frictions that have become standard in quantitative macroeconomics. To this I add the interaction from the simple model: heterogeneous banks sell domestic bonds to rationally inattentive households. I estimate the model using standard macroeconomic data and key series from the savings data in the empirical part of the paper.

This quantitative exercise is possible because of the novel theoretical approach developed in the simple model. Existing macroeconomic models with limited shopping around for prices or interest rates (e.g. McKay (2013), Kaplan and Menzio (2016)) mostly have households engaging in costly search following Burdett and Judd (1983), which outside of simple cases are not usually tractable enough to estimate. I retain many of the qualitative features of the Burdett-Judd model⁵, while keeping the model sufficiently tractable that the interaction of households and banks can be embedded into a quantitative DSGE model, and solved and estimated using standard techniques.

I find that variable attention amplifies the consumption impact of most shocks⁶, as in the simple model. This effect is substantial: the consumption response in the estimated model (cumulated over a year) to government spending and TFP shocks is 43% and 28% larger respectively than if attention is held at steady state. These two shocks explain the largest shares of consumption variation. Overall, the variance of consumption is 17% larger in the baseline model than if variable attention is shut off in this way.

The presence of this amplification has an important policy implication. The extra volatility due to variable attention can be substantially reduced if the marginal cost of information is reduced. Halving the cost of information reduces the variance of consumption by 10%. Policies aimed at providing households with information and facilitating easy product comparisons in this market could therefore lead to lower business cycle volatility.

An additional implication of countercyclical attention is that it can explain a substantial portion of the risk premium shocks typically found to be important in estimated macroeconomic models. Changes in attention affect the model in the same way as risk premium shocks: they change the interest rate faced by the household relative to the policy rate from the central bank. The difference is that attention is an endogenous choice variable. It is not that risk premium shocks cause recessions, but that other kinds of contractionary shock cause attention to rise. An estimation of the model without information costs finds that risk premium shocks explain 25% and 19% of the variance of

⁵The exception is that firms (banks) in my model are not identical, so interest rate heterogeneity is only partly determined by attention, and partly by cost heterogeneity. However, this cost heterogeneity is in fact useful to help the model match the behaviour of interest rate dispersion over the business cycle.

⁶A small number of shocks have large effects of on interest rate dispersion that just outweigh the marginal utility effect discussed above, in which cases attention falls with consumption, weakening the shock slightly. Only one of these accounts for a non-negligible share of consumption and output volatility.

consumption and output respectively. With inattention these shares both fall to 10%. The extra volatility is attributed to supply shocks, notably TFP and price markup shocks, and there is also a larger role given to government spending shocks.

Related Literature. There is a very large literature studying how information frictions affect the business cycle. Many of these papers study frictions in the information agents receive about continuously distributed exogenous shocks (e.g. Mackowiak and Wiederholt, 2015)⁷, or in agents' information about the reaction functions of other agents and the relationships of endogenous variables to shocks (e.g. Eusepi and Preston, 2011)⁸. Unlike these papers, the friction I study is over the discrete choice of which bank to use for saving each period.

Specifically, I draw on Matějka and McKay (2015), who show that an agent facing a discrete choice problem with information costs as in the rational inattention literature (Sims, 2003) will choose information such that choice probabilities resemble those from a multinomial logit model, with a 'twist' that reflects the influence of prior beliefs. This form of inattention is used by Dasgupta and Mondria (2018) to study trade shocks in a model in which countries are inattentive to which country to import each good from, and by Acharya and Wee (2019) to examine information frictions in hiring decisions in a search-based labour market model. I extend the rational inattention literature by showing that information frictions in product choices can have substantial implications for the business cycle.

Another way of modelling the friction in financial product choice would be to use costly search or shopping effort. Coibion, Gorodnichenko and Hong (2015) find that households spend more time and effort shopping for groceries when unemployment rises, echoing my findings that attention to savings product choices rises in contractions. Similarly, the model in Kaplan and Menzio (2016) has unemployed households searching harder for low goods prices, so average search effort rises in recessions. The choice of how much attention to pay to the savings product choice in this paper can be seen as an extension of this literature to financial products, which have particular importance for the business cycle as they influence the intertemporal allocation of consumption.

I also contribute to the literature on the drivers of the business cycle, by showing that countercyclical attention provides a structural interpretation for the risk premium shock that is commonly found to be important in estimated macroeconomics models⁹ (e.g.

⁷In general, agents are assumed to understand the model, so tracking exogenous or endogenous variables are equivalent as agents can perfectly map between them. Other examples of models of this kind include Adam (2007, 2009), Angeletos and La'O (2009, 2020), Auclert et al (2020), Lorenzoni (2010), Mackowiak and Wiederholt (2009), Paciello and Wiederholt (2014), and Reis (2006a, 2006b).

⁸See Eusepi and Preston (2018) for a review of these models.

⁹The lack of a clear structural explanation for this shock is one argument Chari, Kehoe and McGrattan (2009) give against the Smets-Wouters model. Fisher (2015) shows theoretically that it can be interpreted

Smets and Wouters (2007), Christiano, Eichenbaum, and Trabandt (2015)). Attention, however, is not exogenous, but is a response to other variables. Explaining risk premium shocks in this way increases the estimated role of supply shocks in the business cycle.

Several other papers have also studied the role of search or information frictions in financial products. A literature starting with Arrow (1987) finds that information frictions are helpful in explaining wealth inequality, as wealthier households have more incentive to process information about saving and investment choices, and so make better choices and earn higher rates of return on average than less wealthy households¹⁰. In contrast, I focus on a model with a representative household and show that the same information friction amplifies the effect of aggregate shocks on consumption.

Other papers have considered the importance of deposit market frictions for the business cycle. Diebold and Sharpe (1990) and Driscoll and Judson (2013) document significant stickiness in the pass-through from wholesale interest rates to retail deposit rates. Dreschler, Savov and Schnabl (2017) find that this limited pass-through is critical in the transmission of monetary policy, through the effects of policy on bank balance sheets. The mechanism I explore focuses on the effects of deposit frictions on households through their intertemporal consumption decisions, so is a complement to this channel.

Several empirical studies have found evidence of inattention in retail financial markets. Martin-Oliver et al. (2009) and Branzoli (2016) show that interest rate dispersion and the incidence of choice ‘mistakes’ (choosing an unambiguously dominated product) respectively are less common when consumers have more incentive to pay attention to their product choice. Adams et al. (2019) find substantial inattention to savings product choices in a large information-provision experiment with savers at five retail banks in the UK. I contribute to this literature by studying how that inattention varies over the business cycle, and showing the macroeconomic consequences of that variation.

In a closely related paper, Yankov (2018) shows that estimated search costs are substantial in a structural model of the US deposit market based on Burdett and Judd (1983). I use similar data on the menu of interest rates available for the UK, but I extend his work by combining it with data on how households choose within that menu. The more model-free empirical approach this allows also suggests that costly information is important for household decisions over savings products, bolstering Yankov’s argument¹¹. I also differ from Yankov by focusing on the business cycle variation in attention, rather than whether search frictions can explain observed dispersion.

as a shock to the utility of holding safe assets.

¹⁰Campanale (2007), Lei (2019), and Lusardi, Michaud and Mitchell (2017) show that this is quantitatively important for explaining wealth inequality. McKay (2013) studies the importance of the wealth-search link in the welfare consequences of social security privatisation.

¹¹Interestingly, Yankov finds that attention rises when interest rates rise, in contrast to my results. I discuss explanations for this difference in section 4.

The rest of the paper is organised as follows: I develop a partial equilibrium model of rationally inattentive households interacting with heterogeneous banks in section 2. In section 3 I detail the data sources I use, and some institutional background on UK savings markets. I examine this data, showing the dispersion in interest rates and studying household choices within that distribution in section 4. In section 5 I quantify the impact of variable attention on the business cycle by estimating a medium-scale New Keynesian model of the UK incorporating the interaction modelled in section 2. Section 6 concludes.

2 Partial Equilibrium Model

In this section I build a simple partial equilibrium model of rationally inattentive households and heterogeneous banks. Households can pay a utility cost to obtain more information about which of a finite set of banks is offering the best interest rates each period. With more information they will achieve a higher interest rate relative to the distribution of rates on offer. I show that attention (the quantity of information processed) is driven by the marginal utility of income and the dispersion of interest rates. Contractions therefore cause attention to rise, because the marginal utility of income rises. Higher attention makes the deposit market more competitive, so banks increase the rates they offer, further increasing the interest rates households experience. Higher interest rates cause further falls in consumption through the household Euler equation.

2.1 Savings Products

To generate interest rate dispersion in the model, I assume that households buy government bonds through banks, some of whom are more efficient than others. Inefficient (high cost) banks offer lower interest rates than their efficient competitors in equilibrium.

There are N banks. Each period t , each bank n buys bonds from the government and sells them on to individuals, both at price 1. In the next period, the government pays the bank $1 + i_t^{CB}$ per bond bought, and the bank pays $1 + i_t^n$ to the individuals it sold to. Bank n also pays a transaction cost χ_t^n per bond. In this partial equilibrium exercise the policy rate is exogenous, but it is endogenous in the quantitative model in section 5.

Bank n chooses the interest rate they offer to individuals i_t^n to maximise profits, taking into account that their market share¹² will depend on how their interest rate compares with the distribution of rates offered by the other banks. Denoting the probability that

¹²Since all savers are identical, the market share equals the probability a saver chooses that bank.

a saver chooses bank n given a rate distribution as $\Pr(n|\hat{i}_t^n, i_t^{-n})$, the bank problem is:

$$i_t^n = \arg \max_{\hat{i}_t^n} \Pr(n|\hat{i}_t^n, i_t^{-n}) \cdot (i_t^{CB} - \hat{i}_t^n - \chi_t^n) \quad (1)$$

This gives the first order condition¹³:

$$\frac{d}{di_t^n} \Pr(n|i_t^n, i_t^{-n}) \cdot (i_t^{CB} - i_t^n - \chi_t^n) = \Pr(n|i_t^n, i_t^{-n}) \quad (2)$$

Interest rates are dispersed if costs χ_t^n are dispersed. A bank with higher costs will choose lower interest rates, accepting a lower market share to prevent a larger fall in their markup.

2.2 Households

In this simple model, households only choose two things each period: consumption, and how much attention to pay to choosing between the different banks. Attention increases the interest rate the household achieves relative to the distribution of rates on offer.

Specifically, I assume that there is a large representative household composed of many individuals. Each period the household decides how much each individual will consume and save, and how much attention they will pay to the choice of savings products, to maximise expected lifetime utility. As in the Rational Inattention literature, ‘attention’ in this model refers to the amount of information that each individual processes about the banks before deciding which bank to choose for their portion of the household’s saving. All asset income is redistributed among individuals each period, so there is no inequality within the household. The household problem is therefore:

$$\max_{c_t, b_t, i_t^e} \sum_t \beta^t (u(c_t) - \mu \mathcal{I}(i_t^e)) \quad (3)$$

subject to

$$c_t + b_t = b_{t-1}(1 + i_{t-1}^e) + y_t \quad (4)$$

$$\mathcal{I}'(i_t^e) > 0, \mathcal{I}''(i_t^e) < 0 \quad (5)$$

Here c_t is consumption, b_t is real bond holdings, i_t^e is the effective interest rate faced by the household (the average over individuals), and y_t is exogenous income.

The novel element of this problem is the term $\mathcal{I}(i_t^e)$, the amount of attention required for the household to earn an effective interest rate i_t^e on assets bought in period t (which

¹³The market share function derived in section 2.2 is smooth as long as savers have less than full information, and is such that equation 2 is sufficient for profit maximisation.

pay off in $t + 1$), which will be derived below¹⁴. The key properties of this function are expressed in condition 5: if the household pays more attention they will earn a higher rate of interest, but the interest rate gain from more attention diminishes as attention grows. Households choose how much attention to pay by balancing the expected future marginal utility of higher interest income with the costs of attention. I have modelled the costs of attention as a simple additively separable utility cost, with a constant marginal cost μ , as is common in the Rational Inattention literature (see for example Mackowiak and Wiederholt (2009)). This can be thought of as costly cognitive effort. In appendix A.1 I show that time costs and monetary costs lead to the same qualitative conclusions.

In the maximisation I allow the household to directly choose the effective interest rate they face, rather than choosing the amount of attention to pay. These two specifications are equivalent, but the first order condition on the effective interest rate has a more readily interpretable form than a first order condition on attention.

The first order conditions comprise an Euler equation and a first order condition on the effective interest rate:

$$u'(c_t) = \beta \mathbb{E}_t(1 + i_t^e)u'(c_{t+1}) \quad (6)$$

$$\beta b_t \mathbb{E}_t u'(c_{t+1}) = \mu \mathcal{I}'(i_t^e) \quad (7)$$

The first order condition on effective interest rates (equation 7) is crucial in unpicking this model. It shows that households choose attention to equalise the marginal utility of higher asset returns next period with the marginal cost of the attention required to achieve it. The marginal utility of higher asset returns is simply the marginal utility of income in the next period multiplied by the amount saved. Attention therefore rises when consumption is expected to be low, as then the marginal utility of future income rises. It is marginal utility in the following period that matters because the bank choices made in period t only change income when the bonds bought pay out in period $t + 1$.

Attention also rises when the marginal information needed to increase effective interest rates ($\mathcal{I}'(i_t^e)$) falls, as this reduces the marginal cost of increasing asset income. After deriving the relationship between attention and effective interest rates below I show that this marginal cost falls when interest rate dispersion rises, as a small increase in attention leads to a small improvement in the probability of an individual choosing a high interest rate bank. With more dispersed interest rates, this improvement has a larger impact on effective interest rates, and so the extra attention needed to marginally increase i_t^e is low.

The first order condition on effective interest rates also implies that a wealthier household will choose to process more information, and so will experience a higher interest rate.

¹⁴This function will depend on the distribution of interest rates offered by banks, but I have dropped this dependence from the written function to save on notation.

This encourages further saving through the Euler equation¹⁵ (equation 6), but the non-concavity this implies is small enough at plausible parameter values that the first order conditions remain sufficient for utility maximisation (see appendix B.1). I assume that households are net savers (government bonds are in positive net supply), so $b_t > 0$ and the household always chooses to process some information.

I now turn to the derivation of $\mathcal{I}(i_t^e)$, from the decisions of individuals, who face a discrete choice Rational Inattention problem as studied in Matějka and McKay (2015).

Individuals start the period with a prior belief about the probabilities of different states of the world, i.e. interest rate distributions and the positions of banks within that. I assume that individuals share information on returning to the household at the end of the period, so all individuals have the same priors. I denote the probability that an individual chooses bank n if they process no information and rely only on the priors as \mathcal{P}_n . Following Matějka, Steiner and Stewart (2017) I refer to this as the ‘predisposition’ towards bank n .

In general, though, individuals will process some information before choosing a bank. They have access to an infinite set of information about banks. If an individual processed enough of that information before making their bank choice - if they paid enough attention - they would be able to precisely identify the best interest rate in the market and choose it with probability 1. However, because attention is costly, the household chooses to limit the amount of information each individual can process before choosing their bank. Intuitively, each individual could visit every bank in the market and observe their interest rate, and so correctly identify the best product in the market, but doing so requires a great deal of effort and so is prohibitively costly. I further assume that individuals cannot share information within the period.

There are therefore two challenges facing an individual. Using terminology from Matějka and McKay (2015), an individual must decide on an *information strategy* (what kinds of information to process given their limited attention capacity) and an *action strategy* (how to translate that information into a bank choice). Formally, we can write this as the individual choosing the covariance of a noisy signal and the true distribution of banks, subject to a constraint on the amount of information about the bank distribution the signal can contain. The individual then observes a realisation from that noisy signal, updates their beliefs and chooses a bank. The quantity of information embodied in a particular signal structure is defined (following Sims (2003)) as the expected reduction

¹⁵The interaction between attention and wealth implies that the model actually has two steady states, one in which all households are identical and another in which some households are wealthy and attentive, while others remain at the borrowing constraint paying no attention. As the data in section 4 is only informative about average household choices, I study the model with identical households. I study the two-agent steady state in a related model in Macaulay (2020).

in entropy between the prior and posterior beliefs about the rankings of banks in the interest rate distribution after observing a realisation of that signal.

Using lemma 1 from Matějka and McKay (2015), we can leave the belief distributions and signal structures in the background, and rewrite the individual's problem in terms of conditional choice probabilities¹⁶. The individual's maximisation problem becomes:

$$\max_{\Pr(n|s_t)} \mathbb{E}_{s_t} \left(\sum_{n=1}^N i_{s_t}^n \Pr(n|s_t) \right) \text{ subject to} \quad (8)$$

$$\mathcal{I}_t = - \sum_{n=1}^N \mathcal{P}_n \log(\mathcal{P}_n) + \mathbb{E}_{s_t} \sum_{n=1}^N \Pr(n|s_t) \log(\Pr(n|s_t)) \quad (9)$$

The choice variable $\Pr(n|s_t)$ is the probability that the individual chooses bank n given the state of the world is s_t , where s_t summarises the interest rate distribution and the order of banks within it. The individual chooses a decision rule (a set of conditional choice probabilities for each possible ranking of banks s_t) in order to maximise their expected interest rate, as the redistribution of asset income across individuals each period renders them risk neutral over interest rates. They maximise subject to the constraint that $\Pr(n|s_t)$ cannot deviate too far from the predisposition \mathcal{P}_n , the choice probability that would be observed for bank n if the individual had access to no more information than their prior. The more attention the household allows individuals to pay, the more their conditional choice probabilities can deviate from these predispositions, towards the unconstrained choice rule in which $\Pr(n|s_t) = 1$ if bank n offers the highest interest rate in state s_t , and $\Pr(n|s_t) = 0$ otherwise.

Solving the individual's rational inattention problem gives a familiar multinomial logit choice rule:

$$\Pr(n|i_t^n, i_t^{-n}) = \frac{\mathcal{P}_n \exp(\frac{i_t^n}{\lambda_t})}{\sum_{k=1}^N \mathcal{P}_k \exp(\frac{i_t^k}{\lambda_t})} \quad (10)$$

Here I have replaced the notation for a state of the world s_t with the interest rate distribution in time t , made up of the rate offered by bank n and the rates at all of their competitors. The variable λ_t is the lagrange multiplier on the attention constraint 9 in the individual problem, or the shadow value of information. As the household increases attention, holding all else equal the constraint becomes less binding and the shadow value of information falls¹⁷.

¹⁶See Matějka and McKay (2015). Intuitively, it is never optimal to use information processing capacity on two distinct signal realisations that imply the same action, so there is a one-to-one mapping from signal realisations to actions. We can therefore solve the problem by looking at actions (i.e. choice probabilities) rather than explicitly solving for signals.

¹⁷Attention could rise without any change in λ_t if the dispersion of interest rates rises. Then the marginal benefit of information for the individual (in terms of increasing their expected interest rate)

The household decides how much each individual will save before knowing whether they have chosen a bank offering a high or low interest rate. Combined with the income sharing around the household, this means that all individuals save the same amount, and the interest rate the household faces across all of their saving is the expected interest rate achieved by each individual's bank choice. It is this average rate that I refer to as the effective interest rate i_t^e :

$$i_t^e = \sum_{k=1}^N i_t^k \Pr(k|i_t^k, i_t^{-k}) \quad (11)$$

Substituting out for the optimal conditional choice probabilities using equation 10, this becomes:

$$i_t^e = \frac{\sum_{k=1}^N i_t^k \mathcal{P}_k \exp(\frac{i_t^k}{\lambda_t})}{\sum_{k=1}^N \mathcal{P}_k \exp(\frac{i_t^k}{\lambda_t})} \quad (12)$$

As attention increases (λ_t falls), individuals successfully choose higher interest rate banks with a greater probability, and so the effective rate experienced by the household rises. The effective interest rate is therefore an increasing function of the probability of successfully choosing high interest savings products, and information processing \mathcal{I} increases when individuals are more discriminating between banks. Therefore $\mathcal{I}'(i_t^e) > 0$. Diminishing returns to attention ensure that $\mathcal{I}''(i_t^e) > 0$.

The predispositions are where this Rational Inattention model allows for more flexibility than the search-based models explored in appendix A.1. If there is some reason, aside from the current interest rate, for individuals to be more likely to choose one bank than another, that can simply be incorporated into \mathcal{P}_n . The model can therefore incorporate some banks having more 'brand recognition' than others, and so attracting inattentive individuals with a higher probability. It can also allow for persistence in the ordering of banks, in which case knowledge of past states of the world is informative about the current state, and so affects the prior probability of choosing particular banks. While it is possible to construct a search model with bank-specific variation in the probability of individuals meeting each bank, which would be necessary to account for these situations, models of this type quickly become intractable (see e.g. Menzio and Trachter (2015) for a setting with one large seller and a continuum of identical small sellers).

I study the case of persistence in bank costs, and so in the positions of banks within the interest rate distribution, in appendix C.1. For the modelling in the main body of the paper however I assume for simplicity that the ordering of bank costs, and so interest rates, has no persistence¹⁸, and no banks have more brand power than others. The

risks, so if the shadow price of information stays constant attention will rise.

¹⁸There is in fact very little persistence in bank interest rate rankings for the products studied in section 4 (see appendix C.2), though this may not be true of all assets. The Burdett-Judd models

predispositions then all equal $1/N$, and the conditional choice probabilities and effective interest rate become:

$$\Pr(n|i_t^n, i_t^{-n}) = \frac{\exp(\frac{i_t^n}{\lambda_t})}{\sum_{k=1}^N \exp(\frac{i_t^k}{\lambda_t})} \quad (13)$$

$$i_t^e = \frac{\sum_{k=1}^N i_t^k \exp(\frac{i_t^k}{\lambda_t})}{\sum_{k=1}^N \exp(\frac{i_t^k}{\lambda_t})} \quad (14)$$

2.3 Implications

I now analyse the novel channel in this model. A shock that causes consumption to fall leads to higher attention, and so households face higher interest rates relative to the distribution of offers. That distribution shifts up as the deposit market gets more competitive. Through both channels household effective interest rates rise, which through the Euler equation amplifies the consumption fall.

The key equations are the consumption Euler equation (equation 6), the first order condition on attention (equation 7), and the bank profit maximisation condition (equation 2). The bank condition was left in section 2.1 in terms of the probability of savers choosing each bank. Substituting in the conditional choice probabilities from equation 13 this becomes:

$$(1 - \Pr(n|s_t)) \cdot (i_t^{CB} - i_t^n - \chi_t^n) = \lambda_t \quad (15)$$

Similarly, using these choice probabilities we have that $\mathcal{I}'(i_t^e) = \lambda_t^{-1}$. The household first order condition on attention therefore becomes:

$$\beta b_t \mathbb{E}_t u'(c_{t+1}) = \mu \lambda_t^{-1} \quad (16)$$

This shows formally that attention rises with interest rate dispersion through a lower $\mathcal{I}'(i_t^e)$, as claimed in section 2.2. As interest rates become more dispersed, the shadow value of information λ rises. If consumption and saving are constant, equation 16 implies that the household responds by increasing attention, partially relaxing the information constraint in the individual problem, and so bringing λ back down. As I abstract from movements in asset supply¹⁹ (b_t), the only other cause of attention variation is the expected marginal utility of income in the following period.

If attention rises, then the bank first order condition (equation 15) implies that the distribution of interest rates will shift up (proof in appendix B.2). Intuitively, higher attention means that the demand facing an individual bank becomes more elastic to

common in this literature also have no persistence, as all price-setters follow identical mixed strategies.

¹⁹I treat them as constant here, but in the quantitative model the shocks to μ are isomorphic to shocks to b_t , so it can be thought of as being subject to exogenous shocks in section 5.

changes in that bank's interest rate relative to their competitors, as choice probabilities can depend more on specific realisations of interest rates. With more elastic demand, markups decrease, and so the interest rates offered to households rise relative to the policy rate. Furthermore, each bank wants to increase their interest rates to keep pace with rate rises at their competitors, because interest rates are strategic complements in this market²⁰.

Therefore if a shock causes consumption to fall, households increase attention, which means they experience higher interest rates than if attention had remained constant through two channels. First, the probability of an individual choosing high interest rate banks rises, increasing the effective interest rate relative to the distribution of rates on offer. Second, the increased competition in the deposit market causes banks to increase the interest rates they offer, so the rate distribution shifts up. Through the consumption Euler equation, this encourages households to delay consumption, and so consumption falls by even more than it would have done without an attention change. Variable attention therefore amplifies shocks to consumption, unless the shock also reduces interest rate dispersion so much that attention actually falls. In section 5 I find that this is rare, so on average variable attention amplifies the consumption effect of shocks.

3 Data

To provide evidence on cyclical attention to savings, I combine data from two sources. To observe the choice set facing households, I digitise 14 years (1996-2009) of monthly editions of Moneyfacts, a magazine for UK financial advisers. To observe household choices within that set, I combine this with data on average interest rates earned on newly opened savings products each month from the Bank of England. In this section I explain the nature of these datasets, and provide some institutional background on the specific savings market I study.

3.1 Data Sources

Each month Moneyfacts magazine publishes tables of the interest rates and product characteristics of the vast majority²¹ of saving and credit products on offer from retail financial institutions in the UK. A key advantage of this data is that it reports all observable dimensions of product heterogeneity which are relevant for savers, which means that

²⁰In appendix B.2 I also show that interest rate dispersion falls when attention rises. Qualitatively, a rise in attention therefore has the same effect as a rise in search effort in Burdett and Judd (1983).

²¹The publishers aim to cover the universe of products, but acknowledge that they may occasionally miss a niche product from a small provider. As I focus on average household choices in a common product category (fixed interest rate saving bonds), the data should contain all relevant products.

the interest rate dispersion I find after controlling for all of these characteristics cannot be explained by observable product differentiation. The magazine reports the full set of relevant characteristics because it is designed for household financial advisers: if savers care about a product characteristic then financial advisers need to know about it.

Of all of the saving (and borrowing) products available in the data, I focus on the specific subset of fixed interest rate savings products, for which the product characteristics are simple and easily quantifiable. This enables me to account for product heterogeneity. In contrast, mortgages and other loans, as well as other more complicated savings products, have many more dimensions of product heterogeneity, and many products have their own idiosyncratic features, made evident by the paragraph of notes accompanying each observation in the data. Such idiosyncrasies would make accounting for product differentiation in interest rate dispersion extremely difficult. Furthermore, Moneyfacts only reports the advertised interest rates on products: for savings products the vast majority of households receive this rate, but for loans there is the potential for adjustments based on the risk of the individual borrower, so it is not possible to cleanly identify the choice set for loans in the data²². Further details on fixed interest rate savings products are given in section 3.2.

Household choices within this market are reported in the Quoted Household Interest Rate published by the Bank of England. This gives the average interest rate earned by households each month on a subset of fixed interest rate savings products which are identical along all the dimensions of product heterogeneity identified in Moneyfacts, so it directly relates to a set of products which are identical except for the interest rate, and which can be easily identified in the Moneyfacts data. Importantly, the average interest rate reported is for accounts opened in that month only, not the stock of all active accounts, which would include accounts opened in previous months when interest rates were different.

There are several Quoted Household Interest Rate series available for fixed rate savings products with different combinations of product characteristics. I focus on the series for products with a term of one year, an investment of £5000, and where interest is paid annually, because the Quoted Household Interest Rate series goes back to 1996 for these products, whereas the series for other combinations of features have only been published since 2009. In addition, this is one of the most common combinations of product features in the market, so my results in section 4 are less affected by outliers than would be the case with a more niche combination of product features.

²²A further advantage of studying fixed interest rate savings products is that they are mainstream products throughout the sample. In contrast, ISAs (another commonly used savings product in the UK) were only introduced in 1999.

3.2 Institutional Background

Retail savings products are provided in the UK by conventional banks and building societies, which offer deposit products to fund mortgage lending²³. Deposits at all of the institutions in the data were covered by deposit insurance up to £35,000 throughout the period studied, substantially above the £5,000 investment size of the products considered (I return to the issue of deposit insurance and bank risk in section 4.1.1). The largest four institutions had 74% of the market for current accounts in 2000 (Vickers, 2011)²⁴, and the largest branch networks. The market for savings accounts is much less concentrated, with a Herfindahl-Hirschman Index between 20% and 30% lower than the current account market between 2000 and 2008 (Vickers, 2011).

Fixed interest rate savings products are commonly used in the UK. The FCA (2015) found that 12% of households held these products²⁵, and they accounted for 20% of all cash savings balances in the UK. Indeed, the Moneyfacts data confirms that the market is large, as there are an average of 200 such products available each month in the sample. Even in the subset of products with the characteristics that match the Quoted Household Interest Rate there are an average of 34 products each month, and all of the major banks are present.

In addition to the advantages of studying fixed interest rate bonds described above, there are two other factors which aid analysis of choices in this particular market. Firstly, product bundling is uncommon in this market. In its 2015 report, the FCA found that 76% of savers using fixed rate bonds use an institution which is not their ‘main provider of financial services’. I can observe instances of explicit bundling: if, for example, a particular fixed rate product can only be purchased by someone with a current account at that bank, that is noted in the data. I do not remove the few products for which this is the case before analysing the data because they are not removed in the Quoted Rate data, but removing them does not substantially change the distribution of offered rates.

²³The main differences between building societies and banks are that building societies are owned by their customers, and are more limited than banks in how much of their funding can come from wholesale money markets. I will not distinguish between the two types of provider as industry experts suggest it is not important for consumer choices (Maundrell, 2017). For a review of the differences see BBC (2005). As the degree of wholesale funding could be related to bank risk, I discuss this in section 4.1.1.

²⁴This market share fell gradually to 64% in 2008, then rose to 77% in 2010 due to mergers in which large banks bought failing rivals (Vickers, 2011). Cruickshank (2000) reports that in 1998 the ‘big four’ banks had just 19% of savings accounts by number of accounts, but this does not account for deposit size so may not accurately reflect their market share.

²⁵In the 2006 Wealth and Assets Survey 5% of UK households reported holding these products, and for those that did they constituted 31% of household asset portfolios on average. The discrepancy between this figure and that of the FCA (2015) could stem from many products in this market being advertised as ‘fixed rate savings accounts’ rather than ‘fixed-term investment bonds with fixed interest rates’ as they are called in the WAS. Some households holding a fixed-rate bond may have therefore mistakenly described it as a ‘savings account’ in the WAS.

Second, the interest rate is the key product feature that matters for households in this market. In a survey of holders of fixed interest rate bonds the FCA (2015) found that the interest rate was the most important product feature for the large majority of savers in this market. Savers hold fixed rate savings bonds as assets, not for transactions or any other purposes. This is important for my analysis, as customer service and the convenience of a large branch network are unobservable product features that I cannot easily control for. That these do not matter much to savers means that this is unlikely to explain much of the interest rate dispersion I find in section 4.1. The presence of a local branch is less important for these products than others because they are of a fixed maturity, so the saver does not need to interact with the bank on as regular a basis, as is the case for products with the potential for continual adjustment (FCA, 2015). In addition, in the model in section 2 I assumed that the only reason households may not choose the highest interest rate product is because of limited information, and in section 4.2 I use the interest rate that households achieve relative to the distribution of offers as a measure of the ‘success’ of their decision making, and argue that variation in this could be driven by variable attention. These exercises assume that savers always prefer a higher interest rate product, so unobservable product features do not weigh heavily on the value households get from their saving products.

4 Empirical Results

In this section I explore household choice using the datasets described in section 3. First, I show that there is substantial heterogeneity in interest rates offered by retail banks which cannot be explained by product heterogeneity. Without interest rate dispersion, the choice of one savings product over another would have no impact on the interest rate that households experience. I then construct a summary statistic for the ‘success’ of household choice, which measures the interest rate households actually achieved relative to the distribution of rates on offer that month and is closely related to attention in the model in section 2. I show that on average, households more reliably choose higher interest rate products when the unemployment rate is high and the average level of interest rates is low, which is consistent with the model.

4.1 Interest Rate Dispersion

Each month in the sample, households could achieve a wide range of interest rates by choosing different products from different providers. The median within-period standard deviation and interquartile range of interest rates are 60 and 75 basis points respectively.

The median within-period average interest rate is 521 basis points.

Clearly, some of this dispersion is due to the fact that these products are not all close substitutes. They differ in the length of the bond, the minimum investment required, and the frequency with which interest is paid. In section 4.1.1 I account for this product heterogeneity, and show that in all months observable product differentiation explains no more than 53% of interest rate dispersion, and in the majority of periods it explains much less than that. I also argue that dimensions of unobservable product heterogeneity, such as perceived bank risk, are unlikely to explain much of the remaining dispersion. I then provide evidence that limited attention is a likely cause of the remaining interest rate dispersion in section 4.1.2. This means that many savers could increase their interest income without changing any other characteristics of their savings product by switching to other providers. Increased attention to the choice of savings products would lead to this kind of switching, which is how attention affects the interest rate households experience.

4.1.1 Interest Rate Dispersion is not explained by Product Differentiation

I show that product differentiation cannot fully explain the substantial interest rate dispersion in the market for fixed interest savings products in two ways. First, for each period I regress the available interest rates on all product characteristics reported by Moneyfacts. The adjusted R^2 of this regression never exceeds 0.53, so in every month a maximum of just over half of the variation in interest rates is explained by observable characteristics. The median adjusted R^2 across the sample is just 0.18. Second, I consider the group of products in the Quoted Rate data, which are identical across all product characteristics except interest rates, and I show that even among these products the mean within-period standard deviation of interest rates is 43 basis points (on a mean interest rate of 520 basis points). Importantly for these exercises, Moneyfacts supplies all relevant observable dimensions of product differentiation, so the remaining interest rate dispersion must be driven by some other factor, such as limited attention.

For the first exercise, I run the following regression each month, where X_i contains all of the product characteristics listed by Moneyfacts for product i in that month:

$$ir_i = \alpha + \beta X_i + \varepsilon_i \tag{17}$$

Across all periods, the median adjusted R^2 for this regression is 0.18, and the maximum adjusted R^2 is 0.53. At most, just over half of the variation in interest rates can therefore be explained by observable product characteristics.

The second exercise considers only a set of products which are identical along all dimensions of product heterogeneity recorded by Moneyfacts. If the market is perfectly

competitive, and unobserved product heterogeneity is negligible, the products considered should have the same interest rate. This is not what is observed. Within the set of products which qualify for the Quoted Household Interest Rate data (see section 3 for details) the mean within-month standard deviation of interest rates is 43 basis points, on an average interest rate of 520 basis points. In October 2000, as an example, savers could earn annual rates of return between 450 and 680 basis points at different banks on a product with identical characteristics (the standard deviation of rates that month is 44 basis points, which is the median across the sample). The histogram of these rates is plotted in figure 1. There is substantial interest rate dispersion which cannot be explained by observable product differentiation.




Figure 1: Histogram of annual interest rates on fixed interest rate bonds and term accounts on offer in October 2000

These two exercises, however, only control for observable product heterogeneity. The dispersion could be due to unobserved heterogeneity of the products or providers that is known to households. One possible source of such unobserved product differentiation is ‘implicit bundling’: if households have a preference for saving with the same institution they use for their current account, mortgage, and other financial services, then providers with a smaller range of offerings in other product areas may have to pay higher interest rates to compensate savers for the lack of this convenience. The evidence collected by the regulator on this market suggests that this is not a substantial driver of interest rate heterogeneity (see section 3.2)²⁶.

²⁶In fact, the mechanism studied in sections 2 and 5 could be reinterpreted as households choosing how much convenience to give up in order to achieve higher interest rates, rather than how much costly attention to pay to achieve those higher rates. The intuition and qualitative results would be identical in such a convenience-driven model.

Another potential driver of interest rate dispersion is bank risk. A bank that is more likely to fail might have to offer savers higher interest rates to compensate them for the risk that the savers will lose their deposits. This is unlikely, however, to be a significant driver of rate dispersion in this market. Throughout the sample deposits in the UK are insured up to £35,000 (£50,000 after October 2008) per depositor per provider, which is far above the £5,000 investments I study. This removes the majority of risk to savers of bank failure, though if the insurance was not perfectly credible, or the insurer was expected to be slow in paying back deposits in failed banks, then it does not completely eliminate risk as an explanation for rate dispersion. In addition, Chavaz and Slutzky (2018) find that deposit rates in the UK are on average uncorrelated with a variety of measures of bank risk. As interest rate dispersion is substantial in every month of my sample, this suggests that risk is not the main driver of the dispersion. Chavaz and Slutzky do find that riskier banks offer higher interest rates when they face spikes in household attention (measured by Google searches), primarily during the 2008 financial crisis. This suggests that risk may explain why the dispersion of interest rates among similar products rises during the financial crisis in my data, but their results on the relationship between risk and deposit rates on average imply that other factors must also be at play. This is supported by the fact that including bank fixed effects in regression 17 still leaves the mean and median unexplained within-month standard deviation of interest rates at 39 and 37 basis points respectively²⁷.

There could, of course, still be other sources of unobserved product differentiation which explain the dispersion of interest rates that I have not considered here. I therefore proceed by arguing from the other side, giving evidence that there are substantial costs of information/search in this market, and therefore that limited attention could explain why interest rate dispersion persists in equilibrium.

4.1.2 Limited Attention is a plausible explanation of Interest Rate Dispersion

The presence of costly search, information, or attention has been proposed as an explanation of equilibrium price dispersion in a large number of papers, both theoretical and empirical, starting with Stigler (1961) (see Baye et al. (2006) for a review). The existence

²⁷This is an inferior way of capturing risk than that of Chavaz and Slutzky (2018), who use proprietary time-varying measures of bank risk from the Bank of England. Adding bank fixed effects ignores changes in bank risk over time, and also removes all variation which causes a bank to offer persistently high or low rates, whether that is driven by risk or not. The model implies that banks offer persistently higher (lower) interest rates if they have a disproportionately small (large) weight in household prior beliefs (see appendix C.1). The regression with bank fixed effects should therefore be taken as further suggestive evidence that the Chavaz and Slutzky results apply to the fixed-rate market specifically, as well as to retail deposits in general.

of interest rate dispersion not accounted for by observed product differentiation is not, however, evidence in itself that households are less than fully informed about the savings products available to them. I therefore provide evidence that information costs, which lead to inattention, are in fact important in this market.

The clearest piece of evidence for the role of information costs, which would make households inattentive, comes from the FCA (and their predecessor the FSA), who regulate the market for savings products in the UK. In a study of retail financial services for the regulator, Cook et al (2002) concluded that:

“Shopping around is not cost free since consumers have to spend time and effort. The extent to which consumers shop around the market will depend on the benefits they think they can get and the costs of them doing it.”

Other reports by the regulator (FSA (2000), FCA (2015)) on this market have similarly concluded that households could benefit if they searched harder for their financial products, but that such search is costly.

In addition to the remarks of the regulator, the founding of Moneyfacts, the magazine from which I obtain the savings product data, is itself evidence that information costs are substantial in retail financial markets. Moneyfacts was founded to provide data on savings and credit products to financial advisers, because until that point obtaining this information for product comparisons had been difficult (Moneyfacts (2019))²⁸. This suggests that it is costly (in time, effort or money) for households to obtain this information from elsewhere: the magazine would not have been founded, and would not keep selling subscriptions, if data on the full set of available savings products was easy to find. Since less than 8% of UK households employ financial advisers (Aegon, 2017) the existence of the magazine has not itself removed the information friction behind saver inattention.

The rapid spread of comparison websites covering savings products in the early 2000s supports this evidence. The largest such comparison site in the UK, MoneySupermarket, brought in £105 million of revenue in 2006 because very large numbers of people visited the site²⁹ each month to compare a variety of products, including savings products (Connon, 2007). Savers would not need to visit a comparison website if they were already fully informed about the products on offer. However, as with the founding of Moneyfacts, these websites did not reduce the cost of information to zero. It still takes time and effort to use the websites, to process the information and translate it to choices. Indeed, in 2019 the Financial Times ran an article about one bank’s strategy for attracting depositors

²⁸Similarly, a rival comparison service MoneySupermarket began in 1987 when the founder realised that it was very difficult for brokers to compare available mortgage deals (Hohler, 2007).

²⁹By the end of 2006 they had 4 million users per year (Hohler, 2007). MoneySupermarket earns revenue by charging firms each time a consumer clicks through from the site to that firm’s product (Connon, 2007). High revenues therefore reflect large numbers of households using the comparison site.

titled “How Monzo is banking on customer apathy” (Kelly, 2019), indicating that savers are not fully attentive to their choices despite the availability of comparison websites.

Other authors have also concluded that inattention plays an important role in retail financial product markets. Martin-Oliver et al. (2009) find evidence that there is less interest rate dispersion among Spanish banks in markets where households have a greater incentive to pay attention. Branzoli (2016) finds that fewer consumers make the mistake of choosing a product which is strictly dominated by another product at the same bank when they have a greater incentive to pay attention to their choices. For the UK, Adams et al. (2019) find evidence of substantial inattention to savings product choices in a large randomised controlled trial using savers at five retail financial institutions.

Finally, I will discuss below how the endogenous attention decisions studied in the model in section 2 can explain the time series variation in how households choose from among the set of offered rates.

4.2 Constructing φ : a summary statistic for household choice

In this section I use the Moneyfacts and Bank of England data to study how successful households are at choosing the highest interest rate product in the market each month. To do this I compute for each month the difference between the average interest rate earned by households opening new accounts and a benchmark rate, the average interest rate on offer at the four largest banks. I argue that a saver paying no attention would face this benchmark rate on average, and any increase in the rates savers face above this can be seen as an improvement in their choices. Normalising this difference by the standard deviation of interest rates on offer that month ensures that the measure is not mechanically affected by changes in the dispersion of interest rates, and gives a statistic that while model-free in construction is closely related to attention in the model.

I construct the ‘no-attention’ benchmark interest rate to reflect a probable predisposition (to use the language of section 2) towards larger market players: small ‘challenger’ banks are likely to be discovered only if the saver does some careful research, as they do not have large numbers of physical branches or large advertising budgets. Specifically, I construct the benchmark rate by taking the average interest rate on offer from the ‘big four’ banks³⁰. Throughout the sample period these four banks hold most of the market share in many retail banking markets, including current accounts and mortgages, and have many more branches than other banks (Office of Fair Trading, 2008). Using this as the benchmark interest rate assumes that households paying no attention to their choice of savings product are likely to go to their closest bank branch, or the bank where they

³⁰These are Barclays, HSBC, Lloyds, and Royal Bank of Scotland.

hold a current account. Alternative benchmarks, such as weighting banks by their number of branches or the size of their balance sheets, would be strongly correlated with this simple benchmark because the big four consistently dominate others on these metrics.

Figure 2 shows the histogram of interest rates available in October 2000 on the subset of fixed interest rate savings products which appear in the Quoted Household Interest Data, with the benchmark interest rate shown in red and the quoted rate (the average interest rate achieved on products bought that month) shown in green. The benchmark rate is 106 basis points below the maximum rate that households could achieve³¹. While they do not all get that rate, on average savers do somewhat better than they would have if they paid no attention to their choice, earning an average of 6.24% interest, 50 basis points above the benchmark rate.




Figure 2: Histogram of annual interest rates on fixed interest rate bonds and term accounts on offer in October 2000

The statistic on household choice which I will study is the distance between the household mean and the benchmark rate, normalised by the standard deviation of the interest rate distribution that month. I denote the resulting statistic by φ :

$$\varphi_t := \frac{\mathbf{E}_h i_t - i_t^b}{\sigma(i_t)} \quad (18)$$

φ is a summary statistic on how households chose from among a distribution of interest rates in a given month. Note that φ is homogeneous of degree zero in interest rates, so market-wide trends in the level of nominal interest rates do not mechanically affect φ . If household decisions are driven by real interest rates rather than nominal rates, φ is

³¹The highest rate on offer in October 2000 from a big four bank was 6.1%.

unaffected by changes in inflation expectations for the same reason.

Although this statistic is not derived using any particular model, it is closely related³² to attention in the model in sections 2 and 5. In section 2.2 I showed that when a household pays more attention the effective interest rate they experience rises relative to what they would have achieved if they processed no information and simply followed their predispositions. This corresponds to a rise in the average rate achieved by households relative to the benchmark rate, and so a rise in φ . I also showed that attention is only a function of conditional choice probabilities, so if interest rates all move further apart but choice probabilities stay the same attention has not changed. Normalising the gap between the average achieved rate and the benchmark rate by the standard deviation of interest rates ensures that changes in rate dispersion do not mechanically alter φ .

4.3 Properties of φ

The key innovation of φ relative to existing estimates of information processing in savings markets, aside from not having to rely on any specific structural modelling assumptions, is that it can be measured each month. I can therefore study how choice behaviour changes over time, at a high enough frequency to observe co-movements with aggregate variables over the business cycle. In the graphs below I plot the time series of φ , and show that it is positively correlated with the unemployment rate and negatively correlated with the level of interest rates over the business cycle.



Figure 3: φ over time, 6 month moving averages.

A substantial portion of the variation occurs at business cycle frequencies³³. The

³²In appendix D I show that there is an exact correspondence between φ and attention in the model with two banks, and that attention and φ remain closely related with more banks in the market.

³³Figure 3 plots a moving average of φ to aid visualization. 20% of the sample spectral mass of the

largest falls in φ occur during the growth periods of 2004-2005 and 2006-mid 2008. Shortly after the beginning of the Great Recession in the UK in mid-2008, φ began to rise sharply. There was also a substantial rise in φ from July 2001 - April 2002. Although the UK avoided recession during this period, it was a time of slowing growth, and the unemployment rate rose relative to trend.

These observations suggest a countercyclical pattern in φ , which is confirmed in figure 4. These plots show the (HP-filtered) cyclical component of φ against the cyclical components of the average interest rate in the savings market studied³⁴ and in unemployment. Lower interest rates and higher unemployment are associated with higher φ . These relationships are strongly statistically significant: the slope coefficients on interest rates and unemployment both have p-values below 0.01%.




Figure 4: φ against (unweighted) average interest rates among products considered in the Quoted Household Interest Rate data and unemployment. All series are cyclical components after HP filtering. Black solid lines are from linear regressions, which give $\hat{\varphi} = -0.149\hat{i} + 2.7e-9$, $R^2 = 0.146$ and $\hat{\varphi} = 0.381\hat{u} + 9.3e-10$, $R^2 = 0.170$ respectively. Blue circles are averages of φ and the regressor of interest within groups of observations, grouped by their position within the distribution of the regressor.

unsmoothed series lies between 6 and 50 quarters, the business cycle frequency domain suggested by Beaudry, Galizia and Portier (2020).

³⁴I use the unweighted mean of the interest rates in the market, but all rates tend to move together here, so using the benchmark rate, the Quoted Household Interest Rate, or the interest rate on one year UK treasury bills makes little quantitative difference, and no difference to the qualitative conclusions.

When interest rates are high and unemployment is low, savers choose products with low interest rates, close to those offered by the big four banks. As rates fall and unemployment rises, households move up through the distribution of offered rates, more reliably choosing the higher interest rate products in the market, and so achieving higher interest rates relative to the distribution of offers than they did when average rates were high and unemployment was low³⁵. In appendix E.1 I obtain the same result using alternative versions of φ . In particular, I show that in contractions the average interest rate achieved by households moves closer to the highest interest rate on offer in the market, as well as increasing away from the benchmark rate. I keep to this measure of φ in the main body of the paper, however, because of its close correspondence to attention in the model, which enables it to discipline the role of attention in the model.

These cyclical patterns can be explained by the household attention decisions studied in section 2. In recessions, consumption tends to be low, so the marginal utility of interest income is high, increasing the incentives to pay attention³⁶. In addition, when average rates are low in this market the dispersion of interest rates tends to be high, increasing the benefits of attention³⁷. Finally, if there is a ‘search for yield’ motive, i.e. if there is something about low levels of interest rates that make households want to work harder to increase their returns, this would also encourage greater attention, and so higher φ , when average rates are low³⁸. In the model in sections 2 and 5 I allow for the first two channels to operate, leaving examination of the search for yield mechanism for future work.

I have argued above that bank risk does not play a large role in this market. In appendix E.2 I show that changes in the size and composition of the fixed-rate bond market are also unable to explain the cyclical patterns in φ_t . While other explanations of the data are in principle possible, variable attention is therefore the leading candidate.

These results contrast with those of Yankov (2018), who finds that attention is high when rates are *high*. There are two reasons for this. First, the interest rate dispersion he

³⁵These relationships are not driven by extreme events. The linear relationship between interest rates and φ is not significantly different before and after the run on Northern Rock in September 2007. The relationship between φ and unemployment is significantly steeper before the crisis than after, but it remains positive and strongly significant in both sub-samples.

³⁶Similarly, when unemployment is high the opportunity cost of time spent shopping around is low. This does not feature explicitly in the model as the cost of attention is a simple additively separable utility cost. However, for most shocks consumption and output co-move, and so labour supply rises with consumption, in which case this opportunity cost of time channel is qualitatively the same as the marginal utility of income channel.

³⁷ $Corr(\bar{i}, \sigma(i)) = -0.3$ and is significant at the 0.1% level. This correlation is partly driven by the substantial increase in interest rate dispersion during the crisis, which may be partly due to heightened awareness of bank risk (see section 4.1.1). However, this correlation remains negative and significant if I exclude the crisis periods.

³⁸Search for yield often refers to financial institutions taking on more risk to increase their returns when yields are low (e.g. Martinez-Miera and Repullo, 2017). This is somewhat different from the search for yield mentioned here, as in this setting there is no change in the riskiness of household investments.

documents in US markets is positively correlated with the level of interest rates, which is not the case in my UK data³⁹. This means that incentives to search rise with interest rates in his data, but not in mine. Second, to keep the Burdett-Judd equilibrium tractable enough to estimate he studies a two-period model in which a saver's only source of income is interest on their assets. This means that consumption does not vary much in the model, so the main driver of search is the dispersion of interest rates. To the extent that interest rates and consumption are positively correlated, marginal utility will rise when interest rates fall, implying the benefits of search will rise. In section 5 I find that for the UK data the marginal utility of income is indeed a key driver of information choices.

5 Quantitative Assessment

In this section I study the quantitative significance of cyclical attention to saving in an estimated DSGE model for the UK. Cyclical attention amplifies the consumption response to most shocks, as the marginal utility of income channel described in section 2 is estimated to be powerful. This amplification is substantial, increasing the consumption response to government spending and TFP shocks (the two shocks explaining the largest fraction of consumption variance) by 43% and 28% respectively. Overall, the variance of consumption is 17% higher than if attention is held fixed at steady state.

Changes in attention affect the model in the same way as risk premium shocks, and can in fact explain a substantial portion of the business cycle fluctuations otherwise attributed to the risk premium. The important difference is that attention is an endogenous response to other shocks, and so can be influenced by policy. In particular, the majority of the stabilisation effects of holding attention constant can be achieved by reducing the cost of information μ : reducing μ by 50% reduces the variance of consumption by 10%.

5.1 Model

Since the data in sections 3 and 4 concerns savings markets in the UK, I base the model on the medium-scale DSGE model for the UK of Harrison and Oomen (2010).

5.1.1 Full Information Block

The model is a medium-scale small open economy New Keynesian model, with many of the frictions that have become standard in the quantitative macroeconomics literature.

³⁹This may be because he looks at a range of CD's with different sizes of investment. Estimates of equation 17 show that the size of the investment is strongly related to the rate on offer. If the differences between rates on low and high balances grow when the level of rates rise, this could explain some of the discrepancy.

Households consume domestic goods and imports, monopolistically supply differentiated labour varieties, and save through risk-free domestic and foreign bonds, money, and by investing in capital which they rent to firms. They can vary capital utilisation at a cost. They face external consumption habits, capital adjustment costs, nominal wage adjustment costs (with partial indexation to past wages), and portfolio adjustment costs that introduce a friction in holdings of foreign bonds.

Domestic firms hire utilisation-adjusted capital services and labour to monopolistically produce intermediate goods, which are aggregated by perfectly competitive final goods firms who supply home and export markets. Intermediate goods firms face price adjustment costs with partial indexation to past prices, with different adjustment costs for the home and export markets.

A monetary authority sets the interest rate on domestic government bonds following a Taylor Rule with interest rate persistence. The fiscal authority issues a positive amount of bonds, engages in wasteful government spending, and collects lump sum taxes. With full information the model features Ricardian Equivalence. With rational inattention a debt increase only affects consumption because it increases the incentives to pay attention to savings. Changes in debt are therefore isomorphic to changes in the cost of attention μ (see equation 16), and so without loss of generality I fix the supply of (real) bonds at 1, and allow for shocks to μ . I refer to these as ‘attention shocks’ below, but they could equally be interpreted as shocks to government debt.

Foreign variables (inflation, export demand, relative export prices, interest rates) are assumed to follow a VAR process estimated outside of the model, as in Adolfson et al (2007). Details of this are in appendix F.

There are 11 shocks outside of the information problem: to TFP, government spending, the disutility of labour, the capital adjustment cost, the consumption Euler equation (risk premium shock), the price markup on domestic goods, the nominal interest rate (monetary policy shock), and to each of the four international variables.

The only changes I make to the Harrison and Oomen (2010) model, aside from the introduction of inattention to savings as set out below, are that I use a risk premium shock rather than a discount factor shock, and I change the frequency of the model to monthly to make best use of my monthly attention data⁴⁰. For further details on the model setup please therefore see Harrison and Oomen (2010).

⁴⁰Changing the frequency of the model does not affect the equations, but does mean I adjust some of the calibrated parameters and priors used in the estimation of the model from those used by Harrison and Oomen.

5.1.2 Attention Problem

As in section 2, I assume that the household is made up of many individuals, who each purchase their domestic bonds from one of a finite number of banks. The banks' problem is as in section 2.1. To keep the estimation simple I set the number of banks to 2. The information problem only affects the market for domestic bonds, not foreign bonds or capital.

The only difference this makes to the existing household FOCs is that I replace the nominal policy rate in the consumption Euler equation, the FOC on capital and the money demand equation with the nominal effective interest rate, averaged over individuals in the household. Foreign exchange market participants can buy bonds directly from governments, so the interest rate that matters for UIP is the policy rate.

I assume that each period a ranking of banks is drawn. One bank, which I will refer to as the 'good' bank and index by the superscript g , draws a low cost $\chi_t^g = 0$. The other bank draws a high cost, and so I will refer to them as the 'bad' bank (superscript b). They face $\chi_t^b = \tau_1 + \tau_2(i_t^{CB} - \bar{i}^{CB}) + e^{\zeta_{\tau,t}} > 0$. I allow for this cost to depend on the policy rate as a reduced-form way for the model to capture the observed correlation of interest rate dispersion with the level of policy rates. The mean-zero AR(1) shock $\zeta_{\tau,t}$ causes exogenous fluctuations in interest rate dispersion. There is no persistence in the bank cost rankings: each bank has a 50% probability of drawing the low costs each period.

As in section 2, households choose how much attention individuals pay to choosing between banks. More attention increases the effective interest rate by improving the probability that an individual will choose a high-rate bank, but it comes at an additively-separable utility cost with a constant marginal cost.

The first order condition on attention therefore takes a similar form to equation 16, with the only differences being that μ is now subject to a mean-zero AR(1) shock process $\zeta_{\mu,t}$, the stock of saving is set to 1, and the future marginal utility of income is affected by inflation:

$$\beta \mathbb{E}_t \frac{U'(c_{t+1})}{\Pi_{t+1}} = \mu e^{\zeta_{\mu,t}} \lambda_t^{-1} \quad (19)$$

Each individual faces a discrete choice rational inattention problem over the two banks. Since there is no persistence in the rankings of costs faced by banks, and so in the positions of each bank in the interest rate distribution, individuals have uninformative priors. Solving the rational inattention problem, we therefore have that the probability of choosing bank n given that bank n is the good bank that period is p_t^g :

$$p_t^g = \frac{\exp(\frac{i_t^g}{\lambda_t})}{\exp(\frac{i_t^g}{\lambda_t}) + \exp(\frac{i_t^b}{\lambda_t})} \quad (20)$$

The effective interest rate faced by the household is the average over individuals:

$$i_t^e = p_t^g i_t^g + (1 - p_t^g) i_t^b \quad (21)$$

Banks choose interest rates to maximise expected profits. Their first order condition is the same as equation 15 derived in section 2, which for the good and bad bank respectively reduces to:

$$(1 - p_t^g) \cdot (i_t^{CB} - i_t^g) = \lambda_t \quad (22)$$

$$p_t^g \cdot (i_t^{CB}(1 - \tau_2) - i_t^b - (\tau_1 - \tau_2 \bar{i}^{CB}) - e^{\zeta_{\tau,t}}) = \lambda_t \quad (23)$$

Bank profits and transaction costs are redistributed back to the representative household as a lump sum.

There are therefore 5 new variables not in the Harrison and Oomen (2010) model: i_t^e , λ_t , p_t^g , i_t^g , i_t^b . The new equations are the first order condition on attention (equation 19), the choice probability rule (equation 20), the definition of i_t^e (equation 21), and the two bank first order conditions (equations 22 and 23). There are two new shocks, to attention ($\zeta_{\mu,t}$) and rate dispersion ($\zeta_{\tau,t}$).

5.2 Estimation

I conduct a Bayesian Maximum Likelihood estimation of the model solved to a log-linear approximation. There are 11 standard observable variables: GDP, consumption, inflation, the 3-month treasury bill rate, investment, real wages, hours worked, and foreign inflation, industrial production, interest rates, and relative export prices. The foreign variables are trade-weighted averages of the other G7 countries. On top of these I add 3 observables from the Moneyfacts data: the mean and standard deviation of deposit rates, and the choice statistic φ . I use data from 1993-2009.

The three extra observables require three extra shocks: I have already introduced AR(1) shocks to the cost of attention and the dispersion of bank costs, and for the final shock I allow for i.i.d. measurement error for the mean rate⁴¹.

I follow Harrison and Oomen in setting some parameters to match standard values or long-run features of UK data (details in appendix F). I do the same for τ_1 , the constant in the bank cost function, choosing it to match the steady state dispersion of interest rates in the Moneyfacts data.

For the priors on each variable to be estimated I again follow Harrison and Oomen, adjusting them as required for the change in the frequency of the model. The only new

⁴¹This is preferred by the estimation to including a structural shock to mean interest rates relative to the policy rate. The measurement error does not explain a substantial fraction of the mean rate series in the estimation.

parameters to estimate are the cost of attention μ , the cyclicality of bank costs τ_2 , and the persistence and volatility of the new shocks. μ must be greater than 0, but there are no such restrictions on τ_2 . I choose relatively weak priors for both in the absence of strong evidence for the values they should take⁴².

5.3 Results: Amplification from attention

The key novel parameters in the estimation are the cost of information μ and the cyclicality of bank cost dispersion τ_2 , which are estimated to be 0.0124 and -0.1394, both significantly different from zero, with much tighter estimates than the prior distributions. To interpret these estimates I compare the estimated model to an alternative with the same equations and parameters, but where attention is held at its steady state each period. Switching off cyclical variation in attention in this way substantially weakens the transmission of shocks through the economy, so variable attention amplifies shocks.

In table 1 I report the magnitude of the cumulative response of consumption to each shock over a year in the static attention alternative, relative to the baseline estimated model. A value below 1 implies that consumption responds by less to the shock in the fixed attention model than with variable attention. The shocks are ordered according to the share of consumption volatility they explain, with the most important shock first⁴³.

Table 1: Cumulative consumption response to shocks relative to variable attention baseline.

Shock	Fixed Attention
Govt spending	0.699
TFP	0.783
Markup	1.042
Risk premium	0.949
Foreign demand	0.744
Monetary policy	1.015
Labour disutility	1.046
K adj. cost	0.982
Foreign ir	0.869
Export price	0.447
Foreign inflation	0.495

For most shocks, consumption is substantially less responsive when attention is held at its steady state. For government spending and TFP shocks, which together explain 67% of

⁴²Specifically, I choose a gamma prior for μ , with mean 0.01 and standard deviation 0.5, and a uniform prior for τ_2 with mean 0 and standard deviation 0.5.

⁴³Attention and dispersion shocks are excluded because attention shocks have no effect with fixed attention, and dispersion shocks become isomorphic to risk premium shocks.

consumption volatility in the baseline estimated model, attention variation amplifies the consumption response by 43% and 28% respectively. Overall, the variance of consumption is 17% larger with variable attention than if attention is held at steady state.

The intuition is as in section 2: when a shock causes consumption to fall, the marginal utility of income rises, so attention goes up. More attention increases the effective interest rate within the distribution of offers, and causes that distribution to shift up. The household experiences higher interest rates, and reduces consumption even further.

This is also amplified by a further general equilibrium effect not seen in section 2. After a contractionary shock, variable attention reduces output and inflation relative to where they would be with fixed attention. The monetary authority therefore sets a lower policy rate than with fixed attention. Since τ_2 is estimated to be negative, this lower policy rate leads to greater interest rate dispersion, encouraging even more attention⁴⁴.

Amplification from variable attention remains substantial even though the information problem only applies to a subset of the household portfolio, due to a set of no-arbitrage conditions. For households to hold all types of assets the expected benefits of holding them must all be equal. If the household pays more attention to domestic bonds and so increases their interest rate there, the rate on other assets must adjust to match, and so it does not matter that the information problem does not apply to the whole portfolio. In fact, capital provides an extra channel through which attention amplifies fluctuations: when attention rises the interest rate on domestic bonds exceeds the expected return on capital, so investment drops until the returns are equalised, adding to the contraction.

For three shocks, however, shutting off variable attention leads to larger consumption responses, though only the price markup shock plays a substantial role in consumption fluctuations⁴⁵. This is because interest rate dispersion falls when policy rates rise. If there is a shock that causes a small consumption fall but a large rise in the policy rate, then this dispersion effect will dominate and attention will fall. In this case the interest rates households experience will fall relative to the fixed attention case, mitigating the initial fall in consumption. This dispersion effect is small enough that for most shocks causing consumption and interest rates to move in opposite directions (e.g. G , TFP) the marginal utility of income effect dominates and attention amplifies the shock. However, for markup, monetary policy, and labour disutility shocks there is a large change in inflation, and so in policy rates. Attention therefore comoves positively with consumption, dampening the

⁴⁴The implication that attention is higher when the level of interest rates is lower fits with the empirical results in section 4. It also offers a potential alternative explanation for the finding in Dreschler, Savov and Schnabl (2017) that deposit rates rise relative to treasury bill rates when the level of interest rates falls. They attribute this to market power in deposit markets, but it could also be that households pay more attention when rates fall, causing banks to increase rates on average.

⁴⁵Price markup shocks account for 17.2% of consumption variance in the baseline model, where monetary policy and labour disutility shocks account for 1.2% and 0.9% respectively.

shock, though only by a small amount as the marginal utility effect is strong enough to prevent large changes in attention.

Variable attention therefore amplifies the response of consumption to most shocks. For shocks that cause consumption and output to co-move, such as TFP shocks, then this also amplifies the output response. For other shocks, however, output and consumption move in opposite directions (e.g. government spending shocks), and in those cases the amplification of the consumption effect mitigates the output response to the shock.

5.4 Discussion and Policy Implications

An alternative way to understand the effect of variable attention on consumption is to compare it with a risk premium shock. This shock disturbs a wedge between the interest rate experienced by households and the policy rate, which is precisely the effect of a change in attention to savings⁴⁶. The key difference between attention and risk premium shocks is that attention is an endogenous household choice, so is influenced by policy.

In fact, the correspondence between attention and risk premium shocks means that variable attention can provide a structural explanation of risk premium shocks, which is often absent in DSGE models despite the prominent role for these shocks in fitting such models to the data (see Fisher (2015) for an alternative interpretation). To see the quantitative ability of variable attention to explain risk premium shocks, I compare the baseline estimated model with an otherwise identical model without information frictions. The full information model is estimated in the same way as the baseline, with the same data except for the variables associated with the attention problem.

With no information friction, the risk premium shock explains 25% of the variance of consumption, and 19% of the variance of output⁴⁷. Moving to the baseline model with inattention the risk premium shock becomes substantially less important, explaining 10% of both consumption and output.

Cyclical attention can therefore plausibly explain approximately half of the business cycle volatility otherwise attributed to risk premium shocks in the UK. Very little of the fall in the importance of risk premium shocks is made up for by shocks to attention, which explain just 1% of consumption and output variance in the baseline model. This half of the risk premium shock is therefore mostly explained by an endogenous response

⁴⁶This is as long as the profits and transaction costs of banks are transferred back to the household lump sum. If instead the transaction costs are treated as waste they would enter the goods market clearing condition and so the resource constraint would be affected by changes in attention, while it isn't by risk premium shocks. Since the quantitative exercise finds that transaction costs at banks are very small relative to output this effect is at most small.

⁴⁷It explains the second largest share of consumption variance (after government spending shocks), and the third largest share of output variance (after TFP and foreign demand shocks).

of attention to other shocks. In particular, the share of consumption and output variance explained by TFP and price markup shocks increases when adding the information friction. Government spending also explains a greater share of consumption variation. Full results are in appendix F.

Importantly, endogenous attention choices can be affected by policy, where exogenous risk premium shocks cannot. One policy that has an intuitive effect on attention is to reduce the cost of information, for example through financial education programmes or regulation to ensure clearer disclosure and presentation of bank pricing policies.

After a permanent fall in the cost of information μ , households pay more attention to savings in steady state. This reduces the amplification from variable attention through two channels. First, attention becomes more sharply convex in effective interest rates at higher levels of attention ($\mathcal{I}''(i_t^e)$ increases), and so fluctuations in the marginal utility of income produce smaller fluctuations in attention. Second, greater attention reduces the equilibrium dispersion of interest rates, which reduces the impact of attention fluctuations on effective interest rates. For these reasons, reducing μ by 50% (and keeping all other parameters as in the estimated model) reduces the variance of consumption by 10%.

6 Conclusion

I have presented a novel channel through which aggregate shocks affect consumption. In theory and in data, households are more successful at choosing higher interest rate savings products in contractions, because they pay more attention to their choice when the marginal utility of income is high. An improvement in these savings choices increases the interest rate households face, and so causes current consumption to fall as households postpone more consumption to the future. Countercyclical variation in attention therefore amplifies the consumption response to the shocks that drive the business cycle.

In an estimated model of the UK economy, variable attention amplifies the effect of aggregate shocks on consumption: the variance of consumption is 17% higher than it would be if attention remained constant, and the effect of cyclical attention on some specific shocks is substantially larger than that. Variable attention also explains approximately half of the business cycle fluctuations attributed to risk premium shocks in a full information version of the model.

Since attention, unlike the risk premium shock, is an endogenous choice made by households, it can be affected by policy. In particular, policies aimed at making it easier for households to ‘shop around’ for financial products could reduce business cycle volatility, providing another argument in favour of policies such as financial education and clear disclosure of bank pricing policies.

References

- [1] Acharya, Sushant ; Wee, Shu Lin (2019) Rational Inattention in Hiring Decisions. *American Economic Journal: Macroeconomics*, forthcoming
- [2] Adam, Klaus (2007) Optimal Monetary Policy with Imperfect Common Knowledge, *Journal of Monetary Economics* 54.2 (2007): 267-301.
- [3] Adam, Klaus (2009) Monetary Policy and Aggregate Volatility, *Journal of Monetary Economics* 56.Supp. (2009): S1-S18
- [4] Adams, Paul ; Hunt, Stefan ; Palmer, Christopher ; Zaliauskas, Redis (2019) Testing the Effectiveness of Consumer Financial Disclosure: Experimental Evidence from Savings Accounts. National Bureau of Economic Research Working Paper No. 25718
- [5] Aegon (2017) Only 8% of people speak to an adviser about financial decisions, press release. https://www.aegon.co.uk/content/ukpaw/news/only_8_of_peoplespeaktoanadviseraboutfinancialdecisions.html
- [6] Angeletos, George-Marios ; La'O, Jennifer (2009) Noisy Business Cycles, National Bureau of Economic Research Working Paper No. 14982
- [7] Angeletos, George-Marios ; La'O, Jennifer (2020) Optimal Monetary Policy with Informational Frictions, *The Journal of Political Economy* 128.3 (2020): 1027-064.
- [8] Arrow, Kenneth. J (1987) The Demand for Information and the Distribution of Income, *Probability in the Engineering and Informational Sciences*, 1987, Vol.1(1), pp.3-13
- [9] Auclert, Adrien ; Rognlie, Matthew ; Straub, Ludwig (2020) Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model. National Bureau of Economic Research Working Paper No.26647
- [10] Baye, M. R, Morgan, J, Scholten, P. A. (2006). Persistent Price Dispersion in Online Markets. In Jansen, Dennis (Eds.), (pp. 122-143). Northampton, MA: . Edward Elgar
- [11] BBC (2005) Banks and Building Societies. <http://news.bbc.co.uk/2/hi/business/2965412.stm>
- [12] Beaudry, Paul ; Galizia, Dana ; Portier, Franck (2020) Putting the Cycle Back into Business Cycle Analysis. *American Economic Review* 110.1 (2020): 1-47.
- [13] Branzoli, Nicola (2016) Price dispersion and consumer inattention: evidence from the market of bank accounts. *Working Papers of the Bank of Italy*, N. 1082, 2016

- [14] Burdett, K.; Judd, K. (1983) Equilibrium Price Dispersion, *Econometrica*, 51(4), 955-969
- [15] Campanale, Claudio (2007) Increasing Returns to Savings and Wealth Inequality, *Review of Economic Dynamics*, October 2007, 10 (4), 646-675.
- [16] Chari, V. V. ; Kehoe, Patrick J. ; McGrattan, Ellen R. (2009) New Keynesian Models: Not Yet Useful for Policy Analysis. *American Economic Journal: Macroeconomics* 1.1 (2009): 242-66.
- [17] Chavaz, M. ; Slutzky, P. (2018) Paying Attention to Banks: Evidence from Offshore Deposits. Manuscript.
- [18] Christiano, Lawrence J. ; Eichenbaum, Martin ; Trabandt, Mathias (2015) Understanding the Great Recession. *American Economic Journal: Macroeconomics*, 7, 110–67.
- [19] Coibion, Olivier ; Gorodnichenko, Yuriy ; Hong, Gee Hee (2015) The Cyclicalty of Sales, Regular and Effective Prices: Business Cycle and Policy Implications. *American Economic Review*, Vol.105(3), pp.993-1029
- [20] Connon, Heather (2007) The man who made the tills ring at Moneysupermarket. *The Observer*, London, Sun 22 Jul 2007
- [21] Cook, Malcolm ; Earley, Fionnuala ; Smith, Sarah ; Ketteringham, Jody (2002) Losing interest-how much can consumers save by shopping around for financial products?, FSA Occasional Paper, No. 19, October. Gilbert, R A
- [22] Cruickshank, Don (2000) Competition in UK banking : a report to the Chancellor of the Exchequer. Norwich : Stationery Office
- [23] Dasgupta, Kunal ; Mondria, Jordi (2018) Inattentive importers. *Journal of International Economics*, May 2018, Vol.112, pp.150-165
- [24] Diebold, F.X. ; Sharpe, S.A. (1990) Post-Deregulation Bank-Deposit-Rate Pricing: The Multivariate Dynamics, *Journal of Business and Economic Statistics*, 8(3), pp. 281-291.
- [25] Driscoll, J. C. ; Judson, R. A. (2013) Sticky Deposit Rates: Data and Implications for Models of Price Adjustment, *Federal Reserve Board, Finance and Economics Discussion Series working paper*

- [26] Drechsler, Itamar ; Savov, Alexi ; Schnabl, Philipp (2017) The Deposits Channel of Monetary Policy. *The Quarterly Journal of Economics*, Vol. 132(4), pp.1819-1876
- [27] Eusepi, Stefano ; Preston, Bruce (2011) Expectations, Learning, and Business Cycle Fluctuations. *The American Economic Review*, Oct 2011, Vol.101(6), pp.2844-2872
- [28] Eusepi, Stefano ; Preston, Bruce (2018) The Science of Monetary Policy: An Imperfect Knowledge Perspective, *Journal of Economic Literature* 56.1 (2018): 3-59
- [29] Financial Conduct Authority (2015) Cash savings market study report
- [30] Fisher, Jonas D.M. (2015) On the Structural Interpretation of the Smets–Wouters “Risk Premium” Shock. *Journal Of Money, Credit, And Banking* 47.2/3 (2015): Pp511-516.
- [31] Financial Services Authority (2000) Better informed consumers: Assessing the implications for consumer education of research by BMRB
- [32] Gali, Jordi (2008) Monetary policy, inflation, and the business cycle : an introduction to the new Keynesian framework. Princeton, N.J. : Princeton University Press
- [33] Harrison, Richard ; Oomen, Özlem (2010) Evaluating and estimating a DSGE model for the United Kingdom, *Bank of England Quarterly Bulletin*, First Quarter 2010, Vol.50(1), p.52
- [34] Hohler, Emily (2007) My first million: Simon Nixon of Moneysupermarket. *MoneyWeek*, London. <https://moneyweek.com/31394/my-first-million-simon-nixon-of-moneysupermarket/>
- [35] Kaplan, G. ; Menzio, G. (2016). Shopping Externalities and Self-Fulfilling Unemployment Fluctuations. *The Journal of Political Economy*, 124(3), 771-825.
- [36] Kelly, Jemima (2019) How Monzo is banking on customer apathy, *Financial Times Alphaville*, London. April 23, 2019
- [37] Lei, Xiaowen (2019) Information and Inequality, *Journal of Economic Theory* 184 (2019)
- [38] Lusardi, Annamaria; Michaud, Pierre-Carl ; Mitchell, Olivia S (2017) Optimal financial knowledge and wealth inequality, *Journal of Political Economy*, 2017, 125 (2), 431-477.
- [39] Macaulay, Alistair (2020) The Attention Trap: Rational Inattention, Inequality, and Fiscal Policy. Manuscript.

- [40] Maćkowiak, Bartosz ; Wiederholt, Mirko (2009) Optimal Sticky Prices under Rational Inattention, *American Economic Review*, 99, 769-803
- [41] Maćkowiak, Bartosz ; Wiederholt, Mirko (2015) Business Cycle Dynamics under Rational Inattention *Review of Economic Studies*, 82, 1502-1532.
- [42] Martín-Oliver, Alfredo ; Salas-Fumás, Vicente ; Saurina, Jesús (2009) Informational differentiation, interest rate dispersion and market power. *Applied Economics Letters*, 30 October 2009, Vol.16(16), pp.1645-1649
- [43] Martinez-Miera, David ; Repullo, Rafael (2017) Search for Yield. *Econometrica*, March 2017, Vol.85(2), pp.351-378
- [44] Matějka, Filip ; McKay, Alisdair (2015) Rational inattention to discrete choices: A new foundation for the multinomial logit model, *American Economic Review*, 1 January 2015, Vol.105(1), pp.272-298
- [45] Matějka, Filip ; Steiner, Jakub ; Stewart, Colin (2017) Rational Inattention Dynamics: Inertia and Delay in Decision Making, *Econometrica*, March 2017, Vol.85(2), pp.521-553
- [46] Maundrell, Hannah (2017), quoted in The Independent newspaper: Bank or Building Society? Published 15/02/2017. <https://www.independent.co.uk/money/bank-or-building-society-a7581421.html>
- [47] McKay, Alisdair (2013) Search for financial returns and social security privatization, *Review of Economic Dynamics*, 4/2013, Vol.16(2), pp.253-270
- [48] Menzio, Guido ; Trachter, Nicholas (2015) Equilibrium Price Dispersion with Sequential Search, *Journal of Economic Theory* 160.C (2015): 188-215
- [49] Moneyfacts (2019) Company Overview, Moneyfacts Facebook Page, accessed 28/05/2019
- [50] Office of Fair Trading (2008) *Current account research*, Annexe D of *Personal current accounts in the UK – an OFT market study*
- [51] Paciello, Luigi ; Wiederholt, Mirko (2014) Exogenous Information, Endogenous Information, and Optimal Monetary Policy, *The Review of Economic Studies* 81.1 (2014): 356-388
- [52] Ravn, Morten ; Schmitt-Grohe, Stephanie ; Uribe, Martin (2006) Deep Habits, *Review of Economic Studies*, Vol.73(1)

- [53] Reis, Ricardo (2006a) Inattentive Consumers, *Journal of Monetary Economics* 53.8 (2006): 1761-800
- [54] Reis, Ricardo (2006b) Inattentive Producers, *Review of Economic Studies* 73.3 (2006): 793-821
- [55] Sims, Christopher A. (2003) Implications of Rational Inattention. *Journal of Monetary Economics*, 2003, Vol.50(3), pp.665-690
- [56] Smets, Frank ; Wouters, Rafael (2007) Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach. *American Economic Review*, Vol.97(3), pp.586-606
- [57] Stigler, G. (1961) The Economics of Information, *Journal of Political Economy*, 69(3), 213-225.
- [58] Vickers, J. (2011). Independent Commission on Banking : Final report : Recommendations. London: TSO.
- [59] Yankov, Vladimir (2018) In Search of a Risk-free Asset. Manuscript.

A Alternative Mechanisms

A.1 The core mechanism in alternative models

Here I show that the main mechanism of the inattention model of section 2 is also present in a broad class of models in which households can pay a cost to increase the interest rate they face. This includes a model with frictional search for savings products, as in McKay (2013). For simplicity I assume an exogenously fixed distribution of interest rates.

Consider an infinitely lived household who chooses consumption and saving each period to maximise expected lifetime utility subject to a standard budget constraint, where income comes from an endowment y_t and asset income. Households can choose in period t to pay a cost to increase the interest rate they face i_t^e . That is, to achieve i_t^e they must pay a cost $C(i_t^e)$, where C is an increasing convex function. I will consider two specifications for this cost, one in which the cost is an additively separable cost in the utility function, and another in which it is a monetary cost entering the budget constraint. The utility cost specification could be thought of as time or effort spent searching for products, while the monetary cost would be paying an advisor or intermediary to search on their behalf. The specification in use is determined by the binary variable ϕ : when $\phi = 0$ the cost is

a utility cost, when $\phi = 1$ we are studying the monetary cost specification.

$$\max_{c_t, b_t, i_t^e} \mathbb{E} \sum_{t=0}^{\infty} \beta^t [u(c_t) - (1 - \phi)C(i_t^e)] \quad (24)$$

subject to

$$c_t + b_t + \phi C(i_t^e) = y_t + b_{t-1}(1 + i_{t-1}^e) \quad (25)$$

We obtain a familiar consumption Euler equation, and a first order condition on i_t^e :

$$u'(c_t) = \beta(1 + i_t^e)\mathbb{E}_t u'(c_{t+1}) \quad (26)$$

$$\beta b_t \mathbb{E}_t u'(c_{t+1}) = (1 - \phi)C'(i_t^e) + \phi u'(c_t)C'(i_t^e) \quad (27)$$

The household problem in section 2 is a special case of this problem. The household equates the marginal utility of higher asset income with the marginal cost of achieving such a rise in interest rates. With a diminishing marginal utility of consumption, when expected future consumption falls the marginal utility of higher interest rates rises. If $\phi = 0$ households will respond by paying to increase their interest rate, since C is convex. If $\phi = 1$, households will only pay to increase i_t^e (and so $C'(i_t^e)$) if expected future consumption has fallen *relative to* current consumption, as increasing future asset income is achieved by sacrificing current consumption.

After a persistent contractionary shock, expected future consumption will fall, so households will pay to increase their interest rate⁴⁸, which will cause current consumption to fall further through the consumption Euler equation, amplifying the shock. This is the mechanism explored in section 2: the rational inattention problem is a tractable way to motivate and model the cost $C(i_t^e)$ as a utility cost, and allows for the distribution of interest rates available to be endogenised as a bank pricing equilibrium. It is not, however, the only way to do this. I now show that a model with frictional search for banks also fits into this class of models.

Suppose that the household is made up of many individuals. Many banks offer savings products, with interest rates that are distributed according to some CDF $F(i)$. Individuals can only choose a bank for their saving if they have observed its interest rate. All individuals observe one bank drawn at random from F , then with probability ψ they observe a second bank (again drawn at random) before choosing where to place their savings. The meeting rate ψ is an increasing function of the search effort of the individual,

⁴⁸In the monetary cost specification households will only increase their interest rate if future consumption is expected to fall by more than current consumption. In many business cycle models, including that in section 5, internal persistence gives rise to ‘hump-shaped’ dynamics after shocks, which imply that households would pay to increase rates after a contractionary shock in both cost specifications.

denoted e , which is decided by the household.

If an individual observes the interest rates of two banks, they choose the bank offering the higher interest rate, so the interest rate chosen has distribution $(F(i))^2$. The expected interest rate for an individual before we know how many banks they will observe, that is the effective interest rate faced by the household overall, is therefore:

$$i_t^e = (1 - \psi(e_t)) \int if(i)di + 2\psi(e_t) \int if(i)F(i)di \quad (28)$$

This is increasing in the probability of seeing a second bank $\psi(e_t)$, as the expected maximum of two draws from a distribution must be (weakly) greater than the expectation of a single draw. We can rearrange this to express search effort in terms of the interest rate the household ends up facing:

$$e_t = \psi^{-1} \left(\frac{i_t^e - \int if(i)di}{2 \int if(i)F(i)di - \int if(i)di} \right) \quad (29)$$

The fraction inside the inverse ψ function increases linearly in i_t^e . If there are diminishing returns to effort (ψ is concave) then effort will be a convex function of the desired interest rate. If we think of effort as being (psychologically) costly in its own right, or because it uses up valuable time, then the costs of increasing i_t^e will be a direct cost in the household utility function. As long as there are weakly diminishing returns to effort, and the cost of effort is weakly convex in effort, and at least one of those two curvatures is strict, then we obtain the first specification discussed above: there is a direct cost in utility which is convex in the desired (chosen) level of the interest rate. Formally, if the cost of effort in the utility function is $C_e(e)$, then we have:

$$C(i_t^e) = C_e \left(\psi^{-1} \left(\frac{i_t^e - \int if(i)di}{2 \int if(i)F(i)di - \int if(i)di} \right) \right) \quad (30)$$

$$C''(i_t^e) > 0 \text{ if } C_e''(e_t) \geq 0 \text{ and } \psi''(e_t) \leq 0, \text{ one inequality strict} \quad (31)$$

A.2 Loans and misallocation

Here I discuss two alternative channels through which attention to financial product choice could affect the business cycle: attention to loan choice and misallocation of credit. I argue that they are potentially less powerful than the consumption channel of attention to savings that I study in the main body of the paper.

If attention to both saving and loan choices rises in contractions, then savers will face higher interest rates and so reduce their consumption (the main channel studied in this

paper), but borrowers will on average find out about lower interest rate loans, and so will have an incentive to increase their consumption. Attention to loans may therefore counteract the savings channel, but there are two reasons to expect that attention to loans does not operate in the same way, and is less powerful than attention to savings.

Firstly, the most significant debt for the majority of indebted households is a mortgage, and the evidence from the FCA (2019) suggests that there is strong price competition leading to only limited interest rate dispersion in mortgages. The scope for attention to drive interest rate changes is therefore small, and indeed one reason why this might be the case is that the large sums of money involved lead almost all mortgagors to pay a large amount of attention to their choice of product whatever the state of the economy.

Secondly, it is not clear that attention to loan choice will in fact rise in contractions. For savings, I find that the marginal utility of income is very important in determining the extent of attention, and for savers the marginal utility of income is high in (demand-driven) contractions for two reasons: labour income and asset income are both low, as wages and interest rates are low. In contrast, in such a contraction a debtor sees their labour income fall, but the decline in interest rates leads to lower debt repayments, and so to a greater disposable income. It is not therefore clear that attention to loan choice will rise in contractions: for the most indebted households a fall in interest rates will increase disposable income so much that the marginal utility of income could even fall. Unfortunately, the Moneyfacts data on loans is not suitable for an empirical examination of attention to credit products, as the products tend to be very complicated. The equivalent Bank of England data on quoted household interest rates therefore averages over a set of products with substantially different characteristics, and so the comparison of this with the Moneyfacts panel does not accurately reflect search or attention behaviour.

The second alternative mechanism relates to what banks do with deposits. If higher interest rates reflect more productive investment opportunities for the bank, then as households pay more attention to their savings choices in recessions there will be a reduction in loan misallocation, dampening the output effects of the contraction. This is a very interesting potential mechanism, and I leave detailed study of it to future research.

In the specific case of retail savings, however, it is unlikely that this channel has much effect. Particularly in the UK, retail banks take in deposits in order to fund residential lending. Since there is very little interest rate dispersion in mortgages, it is unlikely that banks offering higher deposit rates are doing so because they have access to more profitable lending opportunities.

This argument does not apply to other forms of saving, such as saving in equities. Greater attention to equity choices in recessions should indeed lead to lower misallocation, which would mitigate the amplification of the business cycle that I find through the

consumption channel. Note, however, that less than 20% of equities in the UK are owned by UK individuals⁴⁹ (ONS 2020), and even that figure masks the fact that many of those individuals hold their equities through managed funds and other institutions. This means that an increase in attention by households can only have a small effect on misallocation, as the majority of equity investment decisions are controlled by professional investors, who should spend all of their time paying as much attention to their choices as possible.

B Proofs

B.1 The household FOCs are sufficient for utility maximisation

Here I prove that the household first order conditions are sufficient for utility maximisation in the simple model (section 2), and in the quantitative model (section 5). First, write the household problem as an unconstrained maximisation by substituting out for consumption using the budget constraint:

$$\max_{b_t, i_t^e, X_t} U = \sum_t \beta^t \left(u \left(\frac{b_{t-1}}{\Pi_t} (1 + i_{t-1}^e) + y_t(X_t) - b_t \right) - \mu \mathcal{I}(i_t^e) + v(X_t) \right) \quad (32)$$

Here I have summarised all choice variables other than saving b_t and the effective interest rate i_t^e in the vector X_t . In the simple model there are no other choice variables, so X_t is empty and non-asset income y_t is exogenous. In the quantitative model X_t includes wage setting, investment in capital and foreign bonds, capital utilisation, and money holdings. Inflation erodes real bond holdings as in the quantitative model. Since it does not feature in the simple model, this proof corresponds to that model if Π_t is set to 1 for all t .

I begin by defining H_s as the Hessian matrix of second-order partial derivatives of this utility function with respect to each choice variable that would result if there was no information friction, and so i_t^e was not a choice variable. The Hessian matrix for the full problem is then:

$$H = \begin{bmatrix} & & & 0 \\ & & & \vdots \\ & H_s & & 0 \\ 0 & \dots & 0 & \frac{\partial^2 U}{\partial b_t \partial i_t^e} \\ & & & \frac{\partial^2 U}{\partial i_t^e{}^2} \end{bmatrix} \quad (33)$$

Here I have used the fact that the only choice variable that i_t^e interacts with in the utility

⁴⁹The ONS does not distinguish between foreign individual and foreign institutional investors, but even with these included it is clear that the majority of equities are not under the direct control of households.

function is b_t . For all other choice variables X_t , $\frac{\partial^2 U}{\partial X_t \partial i_t^e} = 0$. The first order conditions are sufficient for utility maximisation if U is weakly concave, which is true if for any vector x :

$$xHx' = x_s H_s x'_s + 2yz \frac{\partial^2 U}{\partial b_t \partial i_t^e} + z^2 \frac{\partial^2 U}{\partial i_t^{e^2}} \leq 0 \quad (34)$$

Where $x_s = [x_1, \dots, y]$ and $x = [x_s, z]$. If households cannot influence effective interest rates the utility function is concave, as then this is a standard household maximisation problem (identical to that in Harrison and Oomen (2010) in the quantitative model). This implies that $x_s H_s x'_s < 0$.

Assuming a diminishing marginal utility of consumption we have that:

$$\frac{\partial^2 U}{\partial b_t^2} = u''(c_t) + \beta \mathbb{E}_t \frac{u''(c_{t+1})(1 + i_t^e)^2}{\Pi_{t+1}^2} < 0 \quad (35)$$

It is therefore sufficient for the concavity of U to show that for any y, z :

$$y^2 \frac{\partial^2 U}{\partial b_t^2} + 2yz \frac{\partial^2 U}{\partial b_t \partial i_t^e} + z^2 \frac{\partial^2 U}{\partial i_t^{e^2}} \leq 0 \quad (36)$$

Using the definition of U this condition becomes:

$$\begin{aligned} y^2 u''(c_t) + y^2 \beta \mathbb{E}_t u''(c_{t+1}) \frac{(1 + i_t^e)^2}{\Pi_{t+1}^2} + 2yz \beta \mathbb{E}_t u''(c_{t+1}) \frac{(1 + i_t^e) b_t}{\Pi_{t+1}^2} \\ + 2yz \beta \mathbb{E}_t u'(c_{t+1}) \frac{1}{\Pi_{t+1}} - z^2 \mu \mathcal{I}''(i_t^e) + z^2 \beta \mathbb{E}_t u''(c_{t+1}) \frac{b_t^2}{\Pi_{t+1}^2} \leq 0 \end{aligned} \quad (37)$$

The two terms that don't depend on c_{t+1} are both negative by definition. Assuming CRRA utility, so $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$, the remaining terms can be written as:

$$\begin{aligned} \beta \mathbb{E}_t \frac{u''(c_{t+1}) b_t^2}{\Pi_{t+1}^2} \left(\frac{y^2 (1 + i_t^e)^2}{b_t^2} + z^2 - 2yz \left(\frac{c_{t+1} \Pi_{t+1} b_t - \gamma (1 + i_t^e)}{b_t \gamma} \right) \right) \\ = \beta \mathbb{E}_t \frac{u''(c_{t+1}) b_t^2}{\Pi_{t+1}^2} \left(\frac{y^2 c_{t+1} \Pi_{t+1}}{b_t \gamma^2} \left(2\gamma (1 + i_t^e) - c_{t+1} \Pi_{t+1} b_t \right) \right. \\ \left. + \left(z - y \left(\frac{c_{t+1} \Pi_{t+1} b_t - \gamma (1 + i_t^e)}{b_t \gamma} \right) \right)^2 \right) \end{aligned} \quad (38)$$

Since $u''(c_{t+1}) < 0$, U is concave if the terms inside the brackets are positive. The final term is positive by definition. Using the functional form for utility, a sufficient condition

for U to be concave is therefore that:

$$-\frac{\beta y^2 b_t}{\gamma} \mathbb{E}_t \frac{c_{t+1}^{-\gamma}}{\Pi_{t+1}} \left(2\gamma(1 + i_t^e) - c_{t+1} \Pi_{t+1} b_t \right) \leq 0 \quad (39)$$

Therefore the first order conditions are sufficient for utility maximisation as long as consumption, inflation and savings are not too large relative to the coefficient of risk aversion and the effective interest rate.

The qualitative results in section 2 hold as long as this condition is satisfied. In the quantitative model this is easily the case for plausible parameterisations. There $b_t = 1$, and γ is estimated to be significantly greater⁵⁰ than 1. Since steady state consumption and inflation are 0.662 and 1, and steady state $i_t^e = 1/\beta - 1 = 0.003$, in the region of the steady state this condition is comfortably satisfied⁵¹. Consumption and inflation would have to be implausibly high, and interest rates implausibly low, to violate this condition, and indeed the estimation never suggests we approach such a region. The condition for the first order conditions to be sufficient for utility maximisation is therefore weak.

B.2 Equilibrium interest rates rise when attention increases

Here I show that when attention rises, the interest rate distribution shifts up, and rate dispersion falls, just as it does in models based on Burdett and Judd (1983).

First, differentiate the first order condition for bank n (equation 15) with respect to λ_t , denoting $\mathcal{S}_t^n = \frac{\exp(i_t^n/\lambda_t)}{\sum_{k=1}^N \exp(i_t^k/\lambda_t)}$ as the market share of bank n in period t , and $d_t^n = i_t^{CB} - i_t^n - \chi_t^n$ as the profit bank n makes per bond sold:

$$-d_t^n \frac{d\mathcal{S}_t^n}{d\lambda_t} - (1 - \mathcal{S}_t^n) \frac{di_t^n}{d\lambda_t} = 1 \quad (40)$$

Using the definition of \mathcal{S}_t^n :

$$\frac{d\mathcal{S}_t^n}{d\lambda_t} = \frac{\mathcal{S}_t^n(1 - \mathcal{S}_t^n)}{\lambda_t} \frac{di_t^n}{d\lambda_t} - \frac{\mathcal{S}_t^n(1 - \mathcal{S}_t^n)i_t^n}{\lambda_t^2} + \mathcal{S}_t^n \left(\sum_{j \neq n} \frac{\mathcal{S}_t^j}{\lambda_t^2} (i_t^j - \lambda_t \frac{di_t^j}{d\lambda_t}) \right) \quad (41)$$

Substituting this in to equation 40 and rearranging we obtain:

$$\frac{di_t^n}{d\lambda_t} = \frac{1}{\lambda_t(1 - \mathcal{S}_t^n)(\lambda_t + d_t^n \mathcal{S}_t^n)} \left[i_t^n d_t^n \mathcal{S}_t^n (1 - \mathcal{S}_t^n) - \lambda_t^2 - d_t^n \mathcal{S}_t^n \left(\sum_{j \neq n} \mathcal{S}_t^j (i_t^j - \lambda_t \frac{di_t^j}{d\lambda_t}) \right) \right] \quad (42)$$

Finally, from equation 15 we can write $\lambda_t = d_t^n(1 - \mathcal{S}_t^n)$. Using this to substitute out for

⁵⁰It is the inverse of σ^c in table 9, so the posterior mode for γ is 1.44

⁵¹With c_{t+1} and Π_{t+1} at steady state and γ at its posterior mode in the estimation then the term in brackets is positive for all $i_t^e > -0.770$.

d_t^n we obtain:

$$\frac{di_t^n}{d\lambda_t} = \frac{1}{\lambda_t} \left[i_t^n \mathcal{S}_t^n - \lambda_t - \frac{\mathcal{S}_t^n}{1 - \mathcal{S}_t^n} \left(\sum_{j \neq n} \mathcal{S}_t^j (i_t^j - \lambda_t \frac{di_t^j}{d\lambda_t}) \right) \right] \quad (43)$$

The first term inside the square brackets is positive, leading higher attention (lower λ_t) to imply lower interest rates. This comes about because banks with high interest rates see their market share rise when attention rises if all interest rates are constant. This provides an incentive for those banks to decrease interest rates, accept a lower market share and make more profit per bond. For banks with lower interest rates (higher costs) this effect is smaller, and may be outweighed by the second term, which is negative.

There is also an indirect effect on the profit maximising interest rate coming from the behaviour of other banks in the market, which is summarised in the final term of equation 43. Firstly, if there are competitor banks with higher interest rates than bank n , an increase in attention will cause bank n 's market share to drop, encouraging an increase in interest rates to offset this. Secondly, if other banks increase their interest rates when attention rises then $\frac{di_t^j}{d\lambda_t} < 0$, which pushes $\frac{di_t^n}{d\lambda_t}$ down: interest rates are strategic complements. As competitors raise their interest rates bank n becomes less competitive and starts to lose market share, so raises their own interest rate to compete.

For this reason the interest rate rises even at the lowest cost bank in the market for plausible values of attention⁵². This bank gains market share when attention rises, pushing them to cut their interest rate and increase profit per bond, but this is more than outweighed by the incentive to raise rates along with their higher cost competitors. Banks with greater market share and higher interest rates relative to their competitors do however increase their rates by less than those with small shares, which implies that a rise in attention leads to a fall in the dispersion of interest rates, just as a rise in search effort reduces price dispersion in Burdett and Judd (1983).

C Persistent bank costs

C.1 Modelling persistent bank costs

Here I show how persistent bank costs affect equilibrium attention, interest rates, and individual choice probabilities. For simplicity, I keep to the case of $N = 2$ banks, though the intuition holds for greater numbers.

⁵²With the estimated parameters in the quantitative model, all interest rates rise with attention as long as the probability of choosing the highest rate bank remains below 0.92, well above the steady state of 0.56. As the probability an individual identifies the lowest cost bank approaches 1, interest rates approach the Bertrand equilibrium: all banks set price equal to marginal cost except the lowest cost bank which sets their interest rate marginally above the next highest rate and captures all of the market.

Suppose that, as in section 5, each period one bank is ‘good’ (cost χ^g) and the other is ‘bad’ (cost $\chi^b > \chi^g$). There are two possible states of the world: in state 1 bank 1 is good and bank 2 is bad, and in state 2 the ordering is reversed. Unlike in section 5, assume that there is persistence in the state. Specifically, the state of the world, denoted s_t , follows a two-state Markov process, in which $\Pr(s_{t+1} = s | s_t = s) = g$, where $g \geq 0.5$.

Assume that savers know the previous state of the world: they observe whether they chose correctly or not when the interest rate payouts occur⁵³. Their choice problem in period t therefore remains a static problem. The persistence in s_t shows up as a prior belief biased towards the previous period’s realised state, which I assume without loss of generality to be state 1. I drop time subscripts to simplify notation. Savers know the bank policy functions, and so they know what interest rate each bank will set in each state of the world. They therefore face the payoff matrix:

Table 2: Payoff matrix, observed previous state

	s_1	s_2
a_1	$i^{1,1}$	$i^{1,2}$
a_2	$i^{2,1}$	$i^{2,2}$
Prior prob.	g	$1 - g$

Here a_n indicates choosing bank n , and $i^{n,s}$ is the interest rate offered by bank n in state s . This matrix is not, in general, symmetric, because bank policy functions depend on both their costs (i.e. the state of the world) and saver predispositions, so bank 1 will set different interest rates in state 1 than bank 2 would in state 2 if $g \neq 0.5$.

With a marginal cost of information of λ , the probability a saver chooses bank n in state s is as in equation 10:

$$P(n|i^{n,s}, i^{-n,s}, s) = \frac{\mathcal{P}_n^0 \exp(\frac{i^{n,s}}{\lambda})}{\mathcal{P}_n^0 \exp(\frac{i^{n,s}}{\lambda}) + (1 - \mathcal{P}_n^0) \exp(\frac{i^{-n,s}}{\lambda})} \quad (44)$$

The unconditional choice probabilities (predispositions) are found as the solution to two normalization conditions:

$$\frac{\exp(\frac{i^{1,1}}{\lambda})g}{\mathcal{P}_1^0 \exp(\frac{i^{1,1}}{\lambda}) + (1 - \mathcal{P}_1^0) \exp(\frac{i^{2,1}}{\lambda})} + \frac{\exp(\frac{i^{1,2}}{\lambda})(1 - g)}{\mathcal{P}_1^0 \exp(\frac{i^{1,2}}{\lambda}) + (1 - \mathcal{P}_1^0) \exp(\frac{i^{2,2}}{\lambda})} = 1 \quad (45)$$

$$\frac{\exp(\frac{i^{2,1}}{\lambda})g}{\mathcal{P}_1^0 \exp(\frac{i^{1,1}}{\lambda}) + (1 - \mathcal{P}_1^0) \exp(\frac{i^{2,1}}{\lambda})} + \frac{\exp(\frac{i^{2,2}}{\lambda})(1 - g)}{\mathcal{P}_1^0 \exp(\frac{i^{1,2}}{\lambda}) + (1 - \mathcal{P}_1^0) \exp(\frac{i^{2,2}}{\lambda})} = 1 \quad (46)$$

⁵³An exploration of this kind of problem without the assumption that individuals know the history of states (but with exogenous payoffs) can be found in Matějka, Steiner and Stewart (2017).

The \mathcal{P}_1^0 that satisfies these conditions is:

$$\mathcal{P}_1^0 = \frac{e^{\frac{i^{21}}{\lambda}} e^{\frac{i^{22}}{\lambda}} - (1-g)e^{\frac{i^{21}}{\lambda}} e^{\frac{i^{12}}{\lambda}} - g e^{\frac{i^{11}}{\lambda}} e^{\frac{i^{22}}{\lambda}}}{e^{\frac{i^{11}}{\lambda}} e^{\frac{i^{12}}{\lambda}} - e^{\frac{i^{21}}{\lambda}} e^{\frac{i^{12}}{\lambda}} - e^{\frac{i^{11}}{\lambda}} e^{\frac{i^{22}}{\lambda}} + e^{\frac{i^{21}}{\lambda}} e^{\frac{i^{22}}{\lambda}}} \quad (47)$$

Since savers observe past states of the world, their priors are entirely determined by the true previous state and the transition probabilities, neither of which the banks can influence. The bank problem therefore remains static: banks choose interest rates to maximise their instantaneous expected profit, giving the same first order condition as in section 2.1:

$$\frac{d}{di_t^n} P(n|s_t) \cdot (i_t^{CB} - i_t^n - \chi_t^n) = P(n|s_t) \quad (48)$$

I assume that banks take saver predispositions as given when deciding their interest rates. Intuitively, this is similar to the assumptions in the deep habits model of Ravn, Schmitt-Grohe and Uribe (2006), in which consumption habits evolve very slowly over time, so firms have limited ability to influence them in the short run. I take this to the extreme and assume that banks cannot influence predispositions at all in the short run. While predispositions must be consistent with interest rate policies in the long run, banks do not take this into account in their decisions. The bank first order condition is then as in section 2:

$$(1 - P(n|s_t)) \cdot (i_t^{CB} - i_t^n - \chi_t^n) = \lambda_t \quad (49)$$

The only difference is that $\Pr(n|s)$ here includes the predisposition, which comes from the prior beliefs, which are in turn driven by the persistence of bank costs.

To find equilibrium, take equation 49 and equation 44 for each of the four combinations of bank and state, and equation 47 to give 9 equations in 9 variables: the four interest rates, four conditional choice probabilities, and the predisposition towards bank 1. Since this allows \mathcal{P}_1^0 to vary in response to interest rates, this equilibrium can be taken as the steady state of the system after predispositions have had time to adjust.

All of the results from the static cost model still hold: as attention rises dispersion falls and average rates rise. This is true both in steady state when \mathcal{P}_1^0 can adjust and after short-run changes when \mathcal{P}_1^0 is constant. Graphs showing this with some example parameters are in figures 5 and 6 below.


On top of those results, we have two new results⁵⁴. First, increasing the persistence of bank costs reduces the amount of attention savers pay each period, as priors become more informative. This causes bank 1 (which is increasingly likely to be low cost) to offer lower interest rates, as savers will come to them with a high probability anyway. Conversely, bank 2 offers higher rates to try and maintain their market share.

⁵⁴Both results are obtained from a large number of simulations using different parameters.


The second result is that the effective interest rate faced by the large household depends on whether the state of the world is the same as in the previous period or not. Bank 1 is more likely to be the low cost bank, so savers are predisposed to choose them. Bank 1 responds to this predisposition by offering lower interest rates. This only partially offsets the prior belief effect, so savers have $\mathcal{P}_1^0 > 0.5$ in equilibrium. This means that if the state stays the same (bank 1 is low cost), savers are more likely to correctly identify the low cost bank than they are if the state changes. This increases the effective interest rate in state 1. At the same time, interest rates at the low cost bank are lower if that low cost bank is bank 1, as they are reacting to savers predispositions. Average interest rates are therefore higher in state 2, which increases the effective interest rate in state 2 relative to state 1. Which effect dominates depends on the parameter values of the model, but in either case there will be two possible effective interest rates, and whenever there is a transition from one state to the other the effective interest rate will change. In most calibrations the second effect dominates, so effective interest rates are higher in the period immediately after a state transition.

State transitions therefore produce i.i.d. shocks to the household effective interest rate. The shocks are i.i.d. because individuals observe the state at the end of each period. If there was a transition from $s = 1$ to $s = 2$ in period t , in period $t + 1$ prior beliefs would have $\Pr(s = 2) = g$, so the effective interest rate would only be at the transition level for one period before returning to the level associated with an unchanged state - unless there was another transition in period $t + 1$. The probability of being at the transition level of i^e is therefore the probability of a state transition happening, which is the same every period. These shocks are the key qualitative difference between this model and the static cost model in sections 2 and 5.


Figure 5: Long run equilibrium varies with λ




(a) Amount of information processed



(b) Equilibrium interest rates for bank 1 (blue) and bank 2 (red) when they are low (solid) and high (dashed) cost.




(c) Predisposition to choosing bank 1




(d) Effective interest rate if the low cost bank is bank 1 (blue) and bank 2 (red).


Figure 6: Short run equilibrium varies with λ



(a) Amount of information processed



(b) Equilibrium interest rates for bank 1 (blue) and bank 2 (red) when they are low (solid) and high (dashed) cost.



(c) Effective interest rate if the low cost bank is bank 1 (blue) and bank 2 (red).

C.2 Persistence of interest rate rankings in the data

In sections 2 and 5 I assume that the ranking of a bank in the interest rate distribution has no persistence. Table 3 shows the bank transition probabilities between quintiles of the interest rate distribution of the products studied in section 4 over a month and a year. The length of a period in section 5 is one month, but the annual transition probabilities are also relevant since these products have a term of one year, so individual savers buying these products only return to the decision a year later.

Without persistence, every transition probability would equal 0.2. The values on the diagonal of the transition matrices are all greater than this, so there is some persistence in the data. However, the persistence is limited, even in the top and bottom quintiles where it is strongest. If a saver chose a bank in the top quintile of the interest rate distribution in a given period, then a year later when their product matures there is only a 33% probability of that bank still being in the top quintile.

	1	2	3	4	5
1	0.77	0.15	0.04	0.02	0.02
2	0.21	0.47	0.20	0.08	0.04
3	0.04	0.27	0.44	0.19	0.07
4	0.01	0.08	0.30	0.40	0.21
5	0.01	0.03	0.09	0.22	0.65

(a) Monthly

	1	2	3	4	5
1	0.59	0.21	0.11	0.06	0.04
2	0.28	0.25	0.22	0.14	0.11
3	0.17	0.24	0.26	0.19	0.15
4	0.08	0.21	0.26	0.23	0.22
5	0.07	0.15	0.21	0.24	0.33

(b) Annual

Table 3: Bank quintile transition matrices. In each table the cell (n, m) indicates the probability of transitioning from the n th quintile to the m th quintile in the following period.

I test if these transition matrices are significantly different from a matrix where every element is 0.2 (the no-persistence case) with a likelihood ratio test:

$$2 \ln \left(\frac{\prod_{n=1}^5 \prod_{m=1}^5 p_{n,m}}{\prod_{n=1}^5 \prod_{m=1}^5 0.2} \right) \sim \chi_{19}^2 \quad (50)$$

The critical value of the test statistic for 5% significance is 30.1. The monthly and annual transition matrices give test statistics of 31.2 and 7.2 respectively. We therefore cannot reject the hypothesis of no persistence at an annual frequency, and we only marginally reject that hypothesis at the monthly frequency.

D Relationship between attention and φ

D.1 N=2 banks

Here I show that there is an exact correspondence between attention and φ in the model with $N = 2$ banks, as used in section 5. More attention means φ rises because individuals choose the higher interest rate in the market with a greater probability, and so the effective rate faced by the household rises relative to the distribution of rates on offer.

As in section 5, assume priors are uninformative and define p_t^g as the probability an individual chooses the high interest rate bank in period t :

$$p_t^g = \frac{\exp(\frac{i_t^g}{\lambda_t})}{\exp(\frac{i_t^g}{\lambda_t}) + \exp(\frac{i_t^b}{\lambda_t})} \quad (51)$$

Individuals paying no attention to bank choice choose bank n with probability $\mathcal{P}_n = 0.5$, so the benchmark no-attention rate in the model is the unweighted mean of the available interest rates:

$$i_t^{NA} = \mathcal{P}_1 i_t^1 + (1 - \mathcal{P}_1) i_t^2 = 0.5(i_t^1 + i_t^2) \quad (52)$$

In this two-bank case of the model, attention is given by:

$$\mathcal{I}_t = \log(2) + p_t^g \log p_t^g + (1 - p_t^g) \log(1 - p_t^g) \quad (53)$$

Attention is a monotonically increasing function of p_t^g (as $p_t^g \geq 0.5$). The empirical statistic φ is:

$$\varphi_t = \frac{p_t^g i_t^g + (1 - p_t^g) i_t^b - \frac{1}{2}(i_t^g + i_t^b)}{\frac{1}{2}(i_t^g - i_t^b)} \quad (54)$$

This simplifies to:

$$\varphi_t = \frac{p_t^g(i_t^g - i_t^b) - \frac{1}{2}(i_t^g - i_t^b)}{\frac{1}{2}(i_t^g - i_t^b)} = 2p_t^g - 1 \quad (55)$$

That is, in this simple case φ_t is a linear function of the probability an individual successfully chooses the higher interest rate bank. With no information processing, $p_t^g = 0.5$ and so $\varphi_t = 0$. If they process enough information to ensure that they always identify the high interest rate bank, $p_t^g = 1$ and $\varphi_t = 1$. The same intuition holds for the N bank case, though the relationship is no longer so precise. This is shown in appendix D.2.

D.2 N>2 banks

Here I show that φ and attention \mathcal{I} are closely related in the model with N banks and uninformative priors. There are no dynamics to the relationship, so for ease I drop all

time subscripts. Denoting the unweighted mean interest rate (which is again the model's no-attention rate) as \bar{i} , and the standard deviation of interest rates as $\sigma(i)$, the model-implied φ is:

$$\varphi = \frac{\sum_n i^n \Pr(\text{choose } n) - \bar{i}}{\sigma(i)} = \frac{\sum_n i^n \exp(\frac{i^n}{\lambda}) - \bar{i}}{\sum_m \exp(\frac{i^m}{\lambda}) - \bar{i}} \quad (56)$$

First, note that as \mathcal{I} approaches 0, λ tends to infinity, and so when attention is 0, $\varphi = 0$:

$$\lim_{\lambda \rightarrow \infty} \varphi = \frac{\frac{1}{N} \sum_n i^n - \bar{i}}{\sigma(i)} = 0 \quad (57)$$

If attention \mathcal{I} reaches $\log(N)$, then each individual can perfectly identify the highest interest rate bank with probability 1, so denoting this as bank 1 (without loss of generality) we have $\varphi > 0$:

$$\varphi(\mathcal{I} = \log(N)) = \frac{i^1 - \bar{i}}{\sigma(i)} = \frac{\frac{1}{N} \sum_n (i^1 - i^n)}{\sigma(i)} > 0 \quad (58)$$

The information constraint is continuous for $\mathcal{I} \in (0, \log(N))$, so the statements above guarantee that \mathcal{I} and φ are positively related at least in some portions of this range.

To make further progress, consider how φ changes in the model assuming that interest rates are held fixed. The equilibrium interest rate response will be incorporated below. We use the chain rule to write:

$$\frac{d\varphi}{d\mathcal{I}} = \frac{d\varphi}{d\lambda} \frac{d\lambda}{d\mathcal{I}} \quad (59)$$

I start with $\frac{d\lambda}{d\mathcal{I}}$. Substituting the optimal choice probabilities into the information constraint 5 gives:

$$\mathcal{I} = \log(N) + \frac{i^e}{\lambda} - \log\left(\sum_n \exp\left(\frac{i^n}{\lambda}\right)\right) \quad (60)$$

Differentiate this with respect to λ :

$$1 = \frac{d\lambda}{d\mathcal{I}} \left[-\frac{i^e}{\lambda^2} + \frac{1}{\lambda} \frac{di^e}{d\lambda} - \frac{d}{d\lambda} \log\left(\sum_n \exp\left(\frac{i^n}{\lambda}\right)\right) \right] \quad (61)$$

Take the final term in the square brackets:

$$\frac{d}{d\lambda} \log\left(\sum_n \exp\left(\frac{i^n}{\lambda}\right)\right) = -\frac{\sum_n i^n \exp(\frac{i^n}{\lambda})}{\lambda^2 \sum_m \exp(\frac{i^m}{\lambda})} = -\frac{i^e}{\lambda^2} \quad (62)$$

Substituting this back into equation 61 the first and third terms in the square brackets cancel, giving:

$$1 = \frac{d\lambda}{d\mathcal{I}} \frac{1}{\lambda} \frac{di^e}{d\lambda} \quad (63)$$

The sign of $\frac{d\lambda}{d\mathcal{I}}$ is therefore the same as the sign of $\frac{di^e}{d\lambda}$:

$$\frac{di^e}{d\lambda} = \frac{(\sum_n i^n \exp(\frac{i^n}{\lambda}))^2 - (\sum_n \exp(i^{n^2} \frac{i^n}{\lambda}))(\sum_m \exp(\frac{i^m}{\lambda}))}{\lambda^2 (\sum_n \exp(\frac{i^n}{\lambda}))^2} \quad (64)$$

The denominator of this fraction is always positive, so the sign is determined by the sign of the numerator, which after expanding the terms in brackets is:

$$\begin{aligned} \sum_n i^{n^2} \exp(\frac{2i^n}{\lambda}) + \sum_{m \neq n} i^n i^m \exp(\frac{i^n + i^m}{\lambda}) - \sum_n i^{n^2} \exp(\frac{2i^n}{\lambda}) - \sum_{m \neq n} i^{n^2} \exp(\frac{i^n + i^m}{\lambda}) \\ = - \sum_{m \neq n} (i^{n^2} - i^n i^m) \exp(\frac{i^n + i^m}{\lambda}) \end{aligned} \quad (65)$$

Inside the sum, each pair of banks $\{j, k\}$ appear twice: when $m = k, n = j$ and when $m = j, n = k$. For each distinct pair of banks $\{j, k\}$, the terms inside the sum are equal to:

$$\exp(\frac{i^j + i^k}{\lambda})(i^{j^2} - i^j i^k + i^{k^2} - i^k i^j) = \exp(\frac{i^j + i^k}{\lambda})(i^j - i^k)^2 > 0 \quad (66)$$

Each pair of terms inside the sum in equation 65 is therefore positive, and so $\frac{di^e}{d\lambda}$ is negative. That is, when the shadow cost of information in the individual problem falls, the effective interest rises, if we hold the distribution of interest rates constant.

This implies that $\frac{d\lambda}{d\mathcal{I}}$ is also negative. If attention rises, then holding the distribution of interest rates constant the shadow price of attention falls.

Now consider $\frac{d\varphi}{d\lambda}$. Since we have already shown that $\frac{di^n}{d\lambda} < 0$ we have:

$$\frac{d\varphi}{d\lambda} = \frac{1}{\sigma(i)} \frac{di^n}{d\lambda} < 0 \quad (67)$$

Therefore holding the distribution of interest rates constant, φ monotonically increases with attention.

To include equilibrium interest rate changes, write the full derivative of φ with respect to \mathcal{I} as:

$$\frac{d\varphi}{d\mathcal{I}} = \frac{\partial \varphi}{\partial \mathcal{I}}|_{i^n} + \sum_n \frac{d\varphi}{di^n} \frac{di^n}{d\mathcal{I}} \quad (68)$$

We have shown that the first term is positive. As discussed in section 2.3 and appendix B, $\frac{di^n}{d\mathcal{I}} > 0$ for the range of attention encountered in simulations of the model. However, the sign of $\frac{d\varphi}{di^n}$ is ambiguous. To see why, it is helpful to examine the numerator and denominator of φ separately.

The numerator is $\sum_n \Pr(\text{choose } n) i^n - \frac{1}{N} \sum_n i^n$. This will rise when a given bank's rate i^j rises if $\Pr(\text{choose } j) > \frac{1}{N}$. For all $\mathcal{I} > 0$ this is true for the lowest cost bank, and

is not true for the highest cost bank. For intermediate banks the sign of this term will depend on the rate distribution and the level of attention. The second-lowest cost bank, for example, will be chosen with a probability greater than $\frac{1}{N}$ at low and modest levels of attention, and rises in attention will initially increase this probability further. However, as attention gets very high individuals will accurately distinguish even between the top two interest rates in the market, and so the probability of this bank being chosen will fall below $\frac{1}{N}$, eventually reaching 0 when $\mathcal{I} = \log(N)$.

The derivative of the denominator of φ (i.e. $\sigma(i)$) with respect to any individual rate i^j is:

$$\frac{d\sigma(i)}{di^j} = \frac{1}{N\sigma(i)} \left(i^j - \frac{1}{N} \sum_n i^n \right) \quad (69)$$

The denominator of φ therefore rises with i^j if i^j is above the mean interest rate in the market. At moderate levels of attention, both the numerator and denominator of φ rise with the interest rates of the banks offering the highest rates, and fall when the lower rates in the market rise. The overall effect on φ of a rise in any individual rate is therefore ambiguous without further specification of the levels of attention and costs.

Simulations of the model with $N > 2$ banks suggest that these equilibrium interest rate response terms are not sufficient to outweigh the direct effect discussed above that $\frac{\partial \varphi}{\partial \mathcal{I}}|_{i^n} > 0$, so there is a positive monotonic relationship between attention and φ in the model even outside of the simple case with $N = 2$ banks.

Note, however, that this derivation still holds the policy rate constant. The policy rate affects the standard deviation of interest rates in the quantitative model if $\tau_2 \neq 0$. This means that an increase in attention when policy rates are low will have a different effect on φ than it would if the increase happened when policy rates are high. As there is not a one-to-one mapping from policy rates to attention (attention decisions co-move differently with the policy rate with different shocks), the link from attention to φ would not be one-to-one if there were more than 2 banks in the quantitative model.

E Robustness for section 4

E.1 Alternative measures of φ

Here I present two alternatives to the household choice statistic φ , which corroborate the evidence in section 4.3 that households move up through the distribution of interest rates when unemployment is high and the level of average rates is low.

First, I define a new variable φ_{best} in a similar way to φ , but rather than comparing the average rate achieved by households each month with the rate at the big four banks,

I compare it with the highest interest rate available in the market. Intuitively, rather than comparing choices to a ‘no attention’ benchmark, this compares choices to a full information benchmark.

$$\varphi_{best} = \frac{\mathbf{E}_h i_t - i_t^{best}}{\sigma(i_t)} \quad (70)$$

Second, I define φ_{pct} to be the percentile of the interest rate distribution at which the average interest rate achieved by the household sits. This is even more model-free than φ and φ_{best} , taking no stance on the appropriate benchmark for choices. As with the previous two statistics, it is homogeneous of degree 0. The downside is that it does not consider the shape of the rate distribution either side of the average rate achieved by households.

$$\varphi_{pct} = \Pr(i_t^n < \mathbf{E}_h i_t) \quad (71)$$

When households are more successful at choosing the higher interest rate products in the market, φ_{best} is low and φ_{pct} is high. The pairwise correlations between each of the three statistics on household choice (φ , φ_{best} , φ_{pct}), unemployment and mean interest rates are shown in table 4 below⁵⁵.

When unemployment is high and interest rates are low, φ_{pct} and φ are high, while φ_{best} is low. All correlations are strongly significant. The two alternative measures of household choice success therefore deliver the same qualitative implications as those found in section 4: in contractions households move up within the distribution of interest rates, away from the low rate offered by the big four banks and towards the highest rate in the market.

Table 4: Pairwise contemporaneous correlations

	φ	φ_{best}	φ_{pct}	U	\bar{i}
φ	1				
φ_{best}	-0.707***	1			
φ_{pct}	0.657***	-0.736***	1		
U	0.413***	-0.310***	0.371***	1	
\bar{i}	-0.382***	0.275***	-0.394***	-0.785***	1

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

E.2 Market size cannot explain fluctuations in φ_t

Dreschler, Savov and Schnabl (2017) (DSS) show that when the Federal Funds Rate rises in the US, retail banks increase their deposit spreads and deposits flow out of the retail market. Here I show that such switching out of the deposit market and into other asset

⁵⁵As in section 4, all correlations are between the cyclical components of each variable, extracted with a HP filter.

types cannot explain the cyclicity of φ_t , because the proportion of households who hold fixed interest rate bonds does not vary significantly through the Great Recession.

In principle, the switching identified by DSS could drive my empirical findings. If households differ in their propensity to pay attention to savings, then it could be that when the level of interest rates rises the high-attention households switch out of the retail deposit market. The savers that remain buying fixed-rate savings bonds from banks are the low-attention households, and so the average attention of households in the market falls without any individual household changing their attention.

To explore if this switching is occurring, I study waves 1-3 (2006, 2008, 2010) of the Wealth and Assets Survey (WAS). This survey asks a large number of households about their assets, including whether they hold fixed interest rate savings bonds, and if so how large their deposits are in such products. As the three waves span the Great Recession, if the DSS switching effect is driving the cyclicity of φ_t we should find that the proportion of households who holding products studied in section 4 increases over the recession.

The products considered in section 4 were those available with an investment of £5000. In the majority of cases, this means that the minimum investment was £5000, and the maximum was £9999, as above that level the banks usually offer a different product with a potentially different interest rate. I therefore study the proportion of households who hold fixed-rate savings bonds with balances in this range. Table 5 shows the results from regressing a dummy variable indicating whether a household owns a fixed-rate bond of the appropriate size on the wave of the WAS they are in, with wave 1 (2006) as the baseline. The probability of holding the bond is not significantly different in each wave. This remains true if I widen the range of deposit sizes, which I do as a check since a minority of products in the sample have a minimum investment below £5000, or don't have a corresponding product for balances above £10000.

This is not inconsistent with the mechanism in DSS. Table 6 shows that the proportion of households holding very large balances in fixed-rate bonds increased through the Great Recession (following the same method as table 5 with larger balances). While a minority of products studied in section 4 would allow these higher balances, in general households could get higher interest rates by depositing larger balances in banks offering specific large-balance bonds, so if these households are high-attention types they are unlikely to be buying products in the set studied above.

Finally, if the proportion of savers in the fixed-rate market who are high-attention types increases in recessions, we should expect to see an increase through the Great Recession in the average education of those who hold fixed-rate bonds. Table 7 shows that the year in which they participated in the survey does not significantly correlate with the education of a holder of a fixed-rate bond of the size considered in section 4.

Table 5: Proportion of households holding the relevant fixed-rate bonds does not change significantly over the Great Recession

	(1)	(2)
	Hold bond £5,000-£9,999	Hold bond £2,500-£12,499
Wave=2	-0.000575 (-0.56)	0.000431 (0.28)
Wave=3	-0.00112 (-1.05)	-0.000229 (-0.15)
Constant	0.0129*** (19.85)	0.0275*** (29.08)
Observations	72197	72197

t statistics in parentheses

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

Table 6: Proportion of households holding larger fixed-rate bonds does change significantly over the Great Recession

	(1)	(2)
	Hold bond £25,000-£49,999	Hold bond £50,000+
Wave=2	0.00290** (2.91)	0.00169 (1.84)
Wave=3	0.00262** (2.67)	0.00428*** (4.40)
Constant	0.0114*** (19.14)	0.0114*** (20.23)
Observations	72197	72197

t statistics in parentheses

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

Table 7: Educational attainment conditional on holding a fixed-rate bond with balance between £5,000 and £9,999 does not change significantly over the Great Recession

	(1)	(2)
	Has degree level education or above	Has some educational qualification
Wave=2	0.0466 (1.23)	0.00341 (0.11)
Wave=3	0.0216 (0.54)	-0.0294 (-0.80)
Constant	0.269*** (12.00)	0.809*** (40.34)
Observations	1052	1052

t statistics in parentheses

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

F Estimation Details

F.1 Data Sources and Treatment

There are 11 standard observable variables: domestic (UK) GDP, consumption, inflation, the 3-month treasury bill rate, investment, real wages, hours worked, and foreign inflation, industrial production, interest rates, and relative export prices. The foreign variables are trade-weighted averages of the other G7 countries. On top of these I add 3 observables from the Moneyfacts data: the mean and standard deviation of deposit rates, and φ . I use data from 1993-2009.

I follow Harrison and Oomen to source and detrend the standard observables. See their paper for details of the data series. The only difference to their approach is that I use the treasury bill rate at a monthly frequency, from the same source as them. I extract a log-linear trend from all real variables, and transform inflation and interest rates into monthly rates before de-meaning⁵⁶. For the average and standard deviation of interest rates in Moneyfacts I follow the same procedure used for the treasury bill rate. For φ I take logs and extract a linear trend, as with the other real variables. I choose τ_1 to match the average gap between the highest and the (unweighted) mean interest rate available.

F.2 Foreign VAR

Foreign variables are assumed to follow a VAR process estimated outside of the model, as in Adolfson et al (2007). Denoting the vector of foreign variables as Y_t , the VAR process

⁵⁶Harrison and Oomen take separate means for each 'regime' of the data identified by Benati (2006), but the final regime begins in 1993, which is when my data starts.

is:

$$F_0 Y_t = F_1 Y_{t-1} + F_2 Y_{t-2} + \dots + u_t \quad (72)$$

In Adolfson et al (2007) Y_t consists of foreign inflation, output, and interest rates. They impose that:

$$F_0 = \begin{bmatrix} 1, & 0, & 0 \\ 0, & 1, & 0 \\ -\gamma_\pi, & -\gamma_y, & 1 \end{bmatrix} \quad (73)$$

Output and inflation are assumed to be unaffected by contemporaneous shocks to anything other than themselves, but interest rates respond to both. The VAR is over-identified, and they cannot reject the over-identifying restrictions.

I have one extra variable, relative export prices. I begin with the Adolfson et al (2007) restrictions, and add that inflation and output also do not respond contemporaneously to shocks to relative export prices. Furthermore, I assume that the foreign interest rate doesn't respond contemporaneously to shocks to relative export prices, but that relative export prices can respond contemporaneously to all variables. The idea is that the exchange rate can vary rapidly in response to shocks, and this will affect the relative export price. This gives:

$$F_0 = \begin{bmatrix} 1, & 0, & 0, & 0 \\ 0, & 1, & 0, & 0 \\ -\gamma_\pi, & -\gamma_y, & 1, & 0 \\ -\gamma_\pi^p, & -\gamma_y^p, & \gamma_r^p, & 1 \end{bmatrix} \quad (74)$$

The model is over-identified. We cannot reject the over-identifying restrictions (p-value 0.33). The AIC suggests that 12 lags is optimal here, which fits with the quarterly models with 4 lags in both Adolfson et al and Harrison and Oomen.

F.3 Estimation Results

Table 8 gives descriptions of each estimated parameter. Tables 9 and 10 show the estimation results for the baseline model and the full information model in section 5.4 respectively. The variance decomposition in the full information and inattention models discussed in section 5.4, is shown in figure 7.

Table 8: Description of estimated parameters

Parameter	Description	Prior Distribution
σ^c	Intertemporal elasticity of substitution	$N(0.66, 0.20)$
ψ^{hab}	Consumption habit parameter	$\text{Beta}(0.88, 0.05)$
σ^h	Labour supply elasticity	$N(0.43, 0.11)$
χ^k	Capital adjustment cost constant	$N(2000, 600)$
ϵ^k	Indexation to past capital adjustment in capital adjustment cost	$\text{Beta}(0.79, 0.1)$
σ^z	Capital utilization cost elasticity	$N(0.56, 0.168)$
χ^{hv}	Domestic goods price adjustment cost	$N(3268, 980.4)$
ϵ^{hv}	Domestic goods inflation indexation	$\text{Beta}(0.64, 0.1)$
χ^{xv}	Export goods price adjustment cost	$N(526.5, 12.5)$
ϵ^{xv}	Export goods inflation indexation	$\text{Beta}(0.52, 0.05)$
ψ^{pm}	Imported goods Calvo parameter	$\text{Beta}(0.80, 0.15)$
ϵ^m	Imported goods inflation indexation	$\text{Beta}(0.55, 0.05)$
ψ^w	Wage Calvo parameter	$\text{Beta}(0.74, 0.05)$
ϵ^w	Wage inflation indexation	$\text{Beta}(0.83, 0.1)$
θ^p	Taylor Rule inflation weight	$N(0.87, 0.1305)$
θ^y	Taylor Rule output weight	$N(0.11, 0.0275)$
θ^{rg}	Taylor Rule persistence	$\text{Beta}(0.95, 0.01)$
μ	Marginal cost of information	$\text{Gamma}(0.01, 0.5)$
τ_2	Elasticity of inefficient bank costs to the policy rate	$U(0, 0.5)$
$\rho_{t fp}$	Persistence of AR(1) TFP shock	$\text{Beta}(0.96, 0.01)$
$\sigma_{t fp}$	s.d. TFP shock	$\text{InvGamma}(0.0056, 2)$
ρ_g	Persistence of AR(1) government spending shock	$\text{Beta}(0.99, 0.01)$
σ_g	s.d. government spending shock	$\text{InvGamma}(0.0085, 2)$
ρ_x	Persistence of AR(1) shock x	$U(0.79, 0.29)$
σ_{κ^h}	s.d. labour disutility shock	$\text{InvGamma}(0.01, 2)$
σ_{rg}	s.d. monetary policy shock	$\text{InvGamma}(0.000707, 2)$
$\sigma_{\zeta^{hb}}$	s.d. price markup shock	$\text{InvGamma}(0.006, 2)$
σ_{ζ^k}	s.d. capital adjustment cost shock	$\text{InvGamma}(0.06, 2)$
σ_y	s.d. shock y	$\text{InvGamma}(0.025, 2)$

$x = \kappa^h, \zeta^c, \zeta^{hb}, \zeta^k, \mu, \tau_2$ refers to the shock to labour disutility, the risk premium, price markups, capital adjustment costs, information costs, and interest rate dispersion.

$y = \zeta^c, \mu, \tau_1, \tau_2$ is a subset of x , plus τ_1 the measurement error in mean rates.

Table 9: Estimated parameters in baseline model

Parameter	Mode	s.d.	Parameter	Mode	s.d.
σ^c	0.694	0.002	ρ_{κ^h}	0.793	0.001
ψ^{hab}	0.944	0.002	ρ_{ζ^c}	0.778	0.001
σ^h	0.467	0.011	$\rho_{\zeta^{hb}}$	0.815	0.001
χ^k	1807.600	12.000	ρ_{ζ^k}	0.740	0.003
ϵ^k	0.826	0.001	ρ_{μ}	0.790	0.001
σ^z	0.580	0.002	ρ_{τ_1}	NA	NA
χ^{hv}	3477.152	9.150	ρ_{τ_2}	0.809	0.001
ϵ^{hv}	0.633	0.001	μ	0.012	0.001
χ^{xv}	526.443	0.002	τ_2	-0.139	0.001
ϵ^{xv}	0.559	0.002	σ_g	0.040	0.002
ψ^{pm}	0.816	0.001	σ_{κ^h}	0.132	0.005
ϵ^m	0.541	0.001	σ_{rg}	0.000	0.000
ψ^w	0.751	0.001	σ_{tfp}	0.005	0.001
ϵ^w	0.854	0.001	σ_{ζ^c}	0.007	0.003
θ^p	0.779	0.003	$\sigma_{\zeta^{hb}}$	0.001	0.000
θ^y	0.112	0.000	σ_{ζ^k}	0.286	0.009
θ^{rg}	0.976	0.002	σ_{μ}	0.066	0.002
ρ_{tfp}	0.983	0.001	σ_{τ_1}	0.475	0.015
ρ_g	0.987	0.001	σ_{τ_2}	0.003	0.005

Table 10: Estimated parameters in full information model

Parameter	Mode	s.d.	Parameter	Mode	s.d.
σ^c	0.829	0.001	ρ_{κ^h}	0.796	0.002
ψ^{hab}	0.899	0.000	ρ_{ζ^c}	0.893	0.001
σ^h	0.434	0.001	$\rho_{\zeta^{hb}}$	0.533	0.001
χ^k	1956.000	0.800	ρ_{ζ^k}	0.592	0.002
ϵ^k	0.713	0.000	ρ_{μ}	NA	NA
σ^z	0.560	0.000	ρ_{τ_1}	NA	NA
χ^{hv}	3161.463	4.575	ρ_{τ_2}	NA	NA
ϵ^{hv}	0.637	0.001	μ	NA	NA
χ^{xv}	525.653	0.027	τ_2	NA	NA
ϵ^{xv}	0.529	0.000	σ_g	0.025	0.002
ψ^{pm}	0.646	0.001	σ_{κ^h}	0.053	0.004
ϵ^m	0.558	0.000	σ_{rg}	0.000	0.000
ψ^w	0.681	0.000	σ_{tfp}	0.005	0.001
ϵ^w	0.796	0.000	σ_{ζ^c}	0.004	0.001
θ^p	0.792	0.001	$\sigma_{\zeta^{hb}}$	0.001	0.000
θ^y	0.100	0.000	σ_{ζ^k}	0.359	0.02
θ^{rg}	0.957	0.000	σ_{μ}	NA	NA
ρ_{tfp}	0.961	0.000	σ_{τ_1}	NA	NA
ρ_g	0.984	0.000	σ_{τ_2}	NA	NA




Figure 7: Percentage of the variance of consumption and output due to each shock in the full information and variable attention models.

The risk premium shock is displayed here in orange. As described in section 5.4, without information frictions the risk premium shock explains 25% of the variance of consumption, more than any other shock except for government spending (39%). It explains the third largest share of output variance (19%), after TFP shocks (25%) and foreign demand shocks (20%).

In moving to the baseline model with inattention the risk premium shock becomes substantially less important, explaining 10% of both consumption and output. This is not picked up by shocks to attention, which explain 1% of the variance of both consumption and output in the baseline inattention model. Rather, TFP and price markup shocks explain greater shares of output variance: with full information they explain 25% and 14%, but with inattention they explain 32% and 30%. For the variance of consumption, government spending, TFP and price markup shocks explain 39%, 18% and 9% with full information, and 47%, 19% and 17% in the baseline model.