

The Attention Trap: Rational Inattention, Inequality, and Fiscal Policy*

Alistair Macaulay[†]

March 25, 2021

Abstract

I show that if it is costly for households to process information about asset returns, a model with ex-ante identical households features persistent inequality. The steady state has a two-agent structure, with inequality maintained by a complementarity between attention and wealth: wealthy households have stronger incentives to pay attention to asset choices, and so earn higher returns than asset-poor households. Fiscal expansions are less powerful in this model than in a standard model with heterogeneous discount factors, because when an expansion causes poor households to start saving, they also increase their attention. They therefore earn higher interest rates, and so save even more, smoothing the windfall from the policy over a longer time period. I provide evidence for this increase in attention using cross-state variation in uncertainty about savings interest rates in the aftermath of the 2017 Tax Cuts and Jobs Act in the US. The effects of fiscal policy therefore depend not just on the existence of inequality, but also on the cause of that inequality.

Keywords: Rational Inattention, Inequality, TANK, Fiscal Policy

*I thank the editor, Florin O. Bilbiie, and two anonymous referees for extremely valuable feedback and guidance. I am grateful to Klaus Adam, Isaac Baley, Paul Beaudry, Martin Ellison, Cosmin Ilut, Filip Matějka, Matthew McKernan, Michael McMahon, Michalis Rousakis, Veronika Selezneva, Christopher Sims, Ctirad Slavík, and seminar participants at the University of Oxford and CERGE-EI for helpful comments and suggestions.

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

1 Introduction

A recent literature has demonstrated that the effects of aggregate shocks are crucially dependent on their distributional consequences (see Kaplan and Violante (2018) for a review). Models in this literature often assume that inequality is driven by exogenous factors,¹ but an empirical literature has highlighted a variety of forces contributing to wealth inequality which plausibly respond to aggregate shocks, with heterogeneity in asset returns a prominent example (e.g. Benhabib et al., 2019; Fagereng et al., 2020). This raises a question: we know that the *existence* of wealth inequality matters for the transmission of aggregate shocks, but does the *source* of that inequality matter? In particular, does incorporating returns heterogeneity affect shock transmission in these models?

In this paper I show that the source of inequality does matter. I develop a two-agent New Keynesian (TANK) model in which wealth inequality is entirely due to returns heterogeneity, through an ‘attention trap’: a complementarity between wealth and asset returns that emerges from optimal household decisions of how much attention to pay to saving choices. Focusing on fiscal policy shocks, often the point of greatest difference between heterogeneous-agent and representative-agent models (Kaplan and Violante, 2018), I show that debt-funded tax cuts have smaller and more persistent effects in this model than in a comparable TANK model where inequality is driven by discount factor heterogeneity (as in Bilbiie et al., 2013).

This difference arises because the attention decisions that drive returns heterogeneity are endogenous to fiscal policy. After a tax cut, low-wealth households increase their saving. This leads them to pay more attention to their choice of savings product, as they now have more to gain from discovering which products are offering the highest interest rates - they have more ‘skin in the game’ (Arrow, 1987). The subsequent increase in the interest rate faced by poor households pushes them to save even more, and so they smooth the tax cut windfall over more periods than in the heterogeneous discount factor model. The marginal propensity to consume (MPC) of poor households is therefore lower in the attention trap model. I provide empirical evidence that attention does indeed respond to fiscal policy in this way using cross-state variation in a survey-based measure of uncertainty about savings interest rates in the aftermath of the 2017 Tax Cuts and Jobs Act (TCJA) in the US.

¹Some households are at their borrowing constraint in these models because they cannot access financial markets (e.g. Bilbiie, 2008), because they are less patient than ‘saver’ households (e.g. Bilbiie et al., 2013), or because they have received negative productivity shocks (e.g. Kaplan et al., 2018). In two-agent models such as Bilbiie (2008) all households have zero wealth, but only some are on their Euler equation. For simplicity I also refer to this heterogeneity as ‘wealth inequality’, as in this paper and the heterogeneous-agent literature these households are unconstrained because of their liquid wealth.

The attention trap model differs from most TANK models in that households are identical ex-ante, with the same preferences and access to asset markets. Building on the representative-agent model in Macaulay (2020), households must decide how much information to process before choosing between different savings products, some of which offer higher interest rates than others. As in the literature on rational inattention (e.g. Sims, 2003), I refer to this information processing as attention. Attention is costly, but with more attention the household is more likely to identify which savings products are offering the best interest rates.

This information problem generates persistent inequality because attention choices depend on wealth. Wealthier households have more to gain from increasing their interest rate because that higher rate operates over a larger pool of assets, and so they pay more attention and earn higher rates (as in Arrow, 1987). If two households have slightly different levels of wealth, the wealthier household will therefore experience higher interest rates, and so will save more, than the less wealthy household. The difference in interest rates acts just like the heterogeneous discount factors in Bilbie et al. (2013), and as in their model it leads to a steady state with two types of households, wealthy households on their Euler equation and poor households against a borrowing constraint. The poor households do not save, even in steady state, because they optimally pay no attention to choosing between saving products and so face very low interest rates.

In fact, the steady state of this model is identical to that in a model where inequality is maintained by heterogeneous discount factors. The two models also give the same response to any shocks that are sufficiently small that all poor households remain at their borrowing constraint. In that case, poor household attention remains at zero and they continue to consume all of their income.

If, however, the fiscal expansion is sufficiently large that poor households temporarily leave the borrowing constraint, the attention trap and heterogeneous discount factor models diverge because the poor households in the attention trap model start to process information. This reduces their MPC, and so the impact of the policy on output.

In addition, this interaction between fiscal policy and attention means that any features of the environment that affect the costs and benefits of attention will also impact the power of fiscal policy, through poor household attention decisions. Lower costs of information about savings choices increase the attention poor households pay once their saving is positive after a fiscal expansion. This further increases their interest rates and reduces their MPCs, and accordingly reduces the impact of the fiscal expansion. Lower taxes on asset income have a similar effect.

Finally, if the interest rates on offer become less dispersed, tax cuts will have even smaller and more persistent effects on aggregate consumption. The highest interest rate

in the market in steady state is unchanged, because it is pinned down by the wealthy household Euler equation. Lower dispersion therefore implies that the interest rate faced by poor households is higher for any given level of attention, so there are stronger incentives for poor households to save. They will start to save after a smaller fiscal stimulus, and for a given stimulus they will save more than if rate dispersion is high. Despite the fact that lower interest rate dispersion reduces the incentive to pay attention to savings choices, the MPC of poor households is therefore lower when interest rate dispersion is low, and fiscal policy will be weaker as a result.

The key mechanism of the paper is therefore that the attention of poor households rises after large fiscal expansions. I provide evidence for this using cross-state heterogeneity in the effects of the 2017 Tax Cuts and Jobs Act. Liu and Williams (2019) find that GDP responds significantly to corporate tax shocks in fewer than 50% of states. In the model, fiscal expansions affect attention through household income, so tax cuts should have a greater impact on poor household attention in states where economic activity responds positively to such cuts. I therefore take the states that Liu and Williams (2019) find respond significantly to corporate tax cuts to be the treatment group in a difference-in-differences model, and estimate the effect of the TCJA on a measure of uncertainty about the interest rate on savings accounts from the Survey of Consumer Expectations.

As predicted by the attention trap model, uncertainty about savings interest rates falls more after the tax cut in the responsive states. This reaction is concentrated in low income households who are not at risk of being unable to pay their bills. This is consistent with the model, where the attention response comes from poor households who start saving on receipt of a tax cut. There is no significant difference in the paths of uncertainty before the tax cut.

The TCJA highlights that large fiscal expansions are not theoretical curiosities, and it is not the only recent example. Fiscal stimulus packages in the US during the Great Recession and Covid-19 pandemic were substantially larger as a portion of GDP than any expansion considered in this paper. In addition, in a less simplified model with a richer wealth distribution, there will always be some households close to lifting off their borrowing constraints. Even small expansions will therefore lead to some households starting to save, and the analysis in this paper suggests that they play an important role in the effects of shocks on consumption.

The results are particularly pertinent in the current environment of interest rates close to the zero lower bound, especially if interest rates remain low for a prolonged period. The lack of room for conventional monetary easing may mean that fiscal policy has to play a greater role in stabilising future recessions (Blanchard and Summers, 2019).² However,

²See the calls from several central bankers for fiscal expansion in the recession caused by the Covid-19

extremely low interest rates may also lead to lower interest rate dispersion as retail banks are unwilling to pass on negative interest rates to households (Eggertsson et al., 2019). The attention trap model suggests that this may weaken the expansionary power of fiscal policy precisely when it is most needed.

Related literature: This paper contributes principally to the literature on the role of household heterogeneity in macroeconomic policy. Several recent papers have shown that heterogeneity in wealth and access to financial markets can have substantial effects on the power of different policies.³ While many of these papers study models with a rich and realistic distribution of wealth, Debortoli and Galí (2018) show that a simpler two-agent (TANK) model based on Bilbiie (2008), with households divided into standard intertemporal optimisers and hand-to-mouth consumers, can match many of the effects of heterogeneity on policy.⁴ The model I build has a similar two-agent structure, but in contrast to Bilbiie (2008) and others in the TANK literature the hand-to-mouth households are at their borrowing constraint because of endogenous attention decisions. To the best of my knowledge, I am the first to incorporate the insight that heterogeneous asset returns can explain a significant fraction of wealth inequality (e.g Benhabib et al., 2019) into a TANK model, and I show that modeling inequality in this way does alter the transmission of fiscal policies compared to a model with exogenously-driven inequality.⁵

The link between attention to saving choices and wealth at the heart of the attention trap was first proposed by Arrow (1987). Since then, a number of papers have assessed the quantitative importance of this mechanism for wealth inequality in rich heterogeneous agent models.⁶ McKay (2013) also uses a model of this kind to quantitatively study the welfare effects of social security privatisation. This paper extends this literature by showing that the Arrow (1987) mechanism delivers a two-agent structure in an otherwise standard New Keynesian model, and that fiscal policy in such a model is less powerful than in an equivalent two-agent model with exogenous wealth heterogeneity.

In support of this mechanism, Fagereng et al. (2020) find in Norwegian data that wealth is correlated with asset returns even within asset classes. In Sweden, Briggs et al.

pandemic (e.g. Powell (2020), Lagarde (2020)).

³For monetary policy, see Gornemann et al. (2016), McKay et al. (2016), Kaplan et al. (2018), among others. For fiscal policy, see Auclert et al. (2018), Hagedorn et al. (2019), Bayer et al. (2020), among others. Werning (2015) and Acharya and Dogra (2020) show that the cyclical risk of idiosyncratic income risk is critical in determining whether household heterogeneity amplifies or dampens shocks. Relatedly, Bilbiie (2019a) shows in an analytic HANK model that amplification occurs only if the ratio of the income of hand-to-mouth households to total income is countercyclical. For empirical evidence on the effects of household heterogeneity on monetary policy, see Almgren et al. (2019) and Holm et al. (2020).

⁴Other papers on TANK models include Galí et al. (2007), Giambattista and Pennings (2017), Bilbiie (2019b), and Cantore and Freund (2021).

⁵This comparison model, and the consideration of the effects of debt-funded tax cuts, is most closely related to Bilbiie et al. (2013), where inequality is driven by heterogeneity in household discount factors.

⁶See e.g. Campanale (2007), Lusardi et al. (2017), Kacperczyk et al. (2019), Lei (2019).

(2020) find that stock market participation increases after lottery wins, and that a model with plausible levels of financial participation costs over-predicts the size of this response. This indicates that other costs, such as the cost of paying attention, are substantial. I contribute to this empirical literature in finding that the Tax Cuts and Jobs Act (TCJA) led to a fall in uncertainty about savings markets among poor households, as predicted by the model. This is the first evidence directly linking information about savings to fiscal policy. In studying this, I also contribute to the literature on the heterogeneity of the effects of tax shocks across states in the US, and the effects of the TCJA specifically.⁷

A closely related paper is Best and Dogra (2019). To my knowledge, theirs is the only other paper in which endogenous returns heterogeneity leads to ex-ante identical households being divided into wealthy and poor in steady state. Their focus, however, is on the stability of different equilibria, so they build a Ramsey model which is sufficiently tractable to solve analytically. Since my focus is on policy and shock transmission, I extend their work by showing that a two-agent steady state exists in an otherwise standard New Keynesian model, and I explore how the implied wealth inequality responds to and affects the transmission of fiscal policies.

I also contribute to the literature on rational inattention in macroeconomics. Beginning with Sims (2003), many papers have explored how costly attention affects macroeconomic models when agents process information about normally distributed random variables.⁸ While this work has been instructive and valuable, Sims (2006) notes that rational inattention also has potential applications in many situations where the assumption of normality is not appropriate. This paper develops one such application, using the discrete choice framework of Matějka and McKay (2015) to embed rational inattention to savings product choices into a New Keynesian model. In this, I build on Macaulay (2020), in which I explore the behaviour of attention to savings over the business cycle in a representative-agent framework. Rational inattention to discrete choices has also been applied to importing (Dasgupta and Mondria, 2018) and hiring decisions (Acharya and Wee, 2020), among other applications.

Finally, it is often argued that fiscal policy must have a greater role in stabilising shocks as the natural rate of interest falls and policy rates become more likely to hit their lower bound (e.g. Blanchard and Summers (2019)). My model highlights a potential limiting force on the power of fiscal policy in such an environment if low real interest rates are associated with low interest rate heterogeneity. Low rate heterogeneity implies that a given fiscal policy will raise poor households further away from their borrowing constraint, which means they smooth the gains from the expansion more over time,

⁷See e.g. Liu and Williams (2019), Occhino (2019), Zidar (2019), Kumar (2020).

⁸See e.g. Adam (2009), Maćkowiak and Wiederholt (2009, 2015), Afrouzi (2020), among others.

reducing the impact of the policy. Fiscal policy will continue to have an effect at the zero lower bound, but the model suggests that it may not be as powerful in a low interest rate environment as implied by models with exogenous wealth heterogeneity.

The rest of the paper is organised as follows: Section 2 sets out the attention trap model and demonstrates how the feedback from wealth to attention leads to a two-agent structure. Section 3 compares the effects of debt-funded tax shocks in this model with the effects in a comparison model with discount factor heterogeneity. Section 4 provides empirical evidence for the link from tax cuts to attention, and Section 5 concludes.

2 Attention trap model

The model is a standard New Keynesian model in many respects, except that there are two savings products available to households each period, with different interest rates. Households must process costly information to discover which offers the better deal. In this way attention (i.e. information processing) will lead households to achieve higher interest rates on their assets (as in Arrow, 1987).

I restrict households to be net savers, so they always prefer higher interest rates. This simplifies the household information problem, but is not necessary for the results. I discuss this further in Section 2.5. The firms and monetary authority are the same as in many standard New Keynesian models, so I leave their problems for Appendix A.1.

2.1 Savings products

Each period, many financial institutions buy one-period government bonds at a price of $(1 + i_t)^{-1}$, then sell them to households at marginal cost. Every period, the institutions draw their marginal costs from an i.i.d. distribution.⁹ With probability g , they have no marginal cost beyond the bond price paid to the government. I refer to the products from these institutions as the ‘good’ savings products. With probability $1 - g$, the financial institution faces an additional marginal cost of $\tau > 0$. The ‘bad’ products from these institutions are therefore sold to individuals at price:

$$\frac{1}{1 + i_t^b} = \frac{1}{1 + i_t} + \tau \tag{1}$$

The ‘bad’ products have a higher price, and so a lower interest rate ($i_t^b < i_t$). As net savers, households prefer the higher interest rates of the good products to the bad.

⁹Assuming i.i.d. costs simplifies the rational inattention problem facing households, as it means that information processed about dealers in one period has no value in future periods. See Steiner et al. (2017) for results on information processing problems in dynamic settings.

2.2 Households

Households are composed of many individuals. The household decides the consumption and labour supply of each individual, and how much attention each individual will pay to the choice of savings products, subject to a no-borrowing constraint. More attention will imply that the household experiences a higher interest rate.

The household sets the (real) quantity of bonds b_t that each individual will buy before knowing which of the two available bond prices they face, and there is perfect consumption sharing around the household. Each individual chooses one savings product.¹⁰ They choose a ‘good’ product with probability q_t , which is a function of information gathered \mathcal{I} . I refer to this information processing as attention, as is common in the rational inattention literature. There is a constant marginal cost μ to increasing \mathcal{I} .

These assumptions imply that the household faces the average bond price:

$$\frac{1}{1 + \tilde{i}_t} = \frac{q_t}{1 + i_t} + \frac{1 - q_t}{1 + i_t^b} \quad (2)$$

Each individual faces a discrete choice rational inattention problem (see Matějka and McKay, 2015). As they always prefer the good savings product, they seek to maximise the success probability q . Using the entropy-based measure of information processing common in the rational inattention literature (e.g. Sims, 2003), we have for $q \in [g, 1)$:¹¹

$$\mathcal{I}_t = q_t \log \left(\frac{q_t}{g} \right) + (1 - q_t) \log \left(\frac{1 - q_t}{1 - g} \right) \quad (3)$$

With no information processing, individuals choose savings products at random and so choose a good product with probability $q_t = g$. As attention increases, the individual processes more information before deciding on a product, and so chooses a good product with a greater probability. The probability of choosing a good product is concave in information processing, so there are diminishing returns to attention.¹²

The household problem is to choose consumption c_t , labour supply n_t , real bond purchases b_t , and the effective interest rate \tilde{i}_t (reflecting attention choice) to solve:

$$\max_{c_t, n_t, b_t, \tilde{i}_t} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{c_t^{1-\sigma}}{1-\sigma} - \frac{n_t^{1+\varphi}}{1+\varphi} - \mu \mathcal{I}(\tilde{i}_t) \right) \quad (4)$$

¹⁰Redistributing consumption around the household implies that individuals are risk neutral with respect to the bond price. There is therefore no incentive for them to spread their savings across multiple banks, so they choose a single bank for their portion of the household’s saving.

¹¹See Appendix B.1 for derivation.

¹²The attention required to choose a good product with probability arbitrarily close to 1 is finite (proof in Appendix B.1). However, $\frac{dq_t}{d\mathcal{I}_t}$ approaches zero as q_t approaches 1, so the returns to attention tend to zero and no household achieves $q_t = 1$.

subject to

$$c_t + \frac{b_t}{1 + \tilde{i}_t} = \frac{b_{t-1}}{\Pi_t} + w_t n_t + \mathcal{D}_t - T_t \quad (5)$$

$$b_t \geq 0 \quad (6)$$

Equation (5) is the budget constraint in real terms. Π_t is gross inflation, w_t is the real wage, \mathcal{D}_t is profits from firms, and T_t is a lump sum tax. The effective interest rate is an increasing function of the probability of successfully choosing a good savings product q_t , and information processing \mathcal{I} is an increasing function of q_t . Therefore $\mathcal{I}'(\tilde{i}_t) > 0$. Diminishing returns to attention ensure that $\mathcal{I}''(\tilde{i}_t) > 0$ (see Appendix B.2).

The first order conditions comprise an Euler equation, a labour supply condition, and a first order condition on attention:

$$c_t^{-\sigma} = \beta(1 + \tilde{i}_t)\mathbb{E}_t\left(\frac{c_{t+1}^{-\sigma}}{\Pi_{t+1}}\right) + (1 + \tilde{i}_t)\lambda_t \quad (7)$$

$$c_t^\sigma n_t^\varphi = w_t \quad (8)$$

$$b_t c_t^{-\sigma} = \mu(1 + \tilde{i}_t)^2 \mathcal{I}'(\tilde{i}_t) \quad (9)$$

λ_t is the Lagrange multiplier on the borrowing constraint, which is positive when the constraint binds and zero otherwise. The first order condition on attention (9) implies that a wealthier household will choose to process more information, and so will experience a higher interest rate. This encourages further saving through the Euler equation (7), but the non-concavity this implies is small enough at plausible parameter values that the first order conditions remain sufficient for utility maximisation (proof in Appendix B.3). This interaction between attention and wealth is the reason this model has a steady state with two types of households, some wealthy and others at the borrowing constraint, despite those households having identical preferences and access to asset markets. In the remaining sections I denote a variable pertaining to wealthy or poor households with a superscript w or p respectively, and let ω be the proportion of wealthy households.¹³

2.3 Government

Government issues a positive amount of short term bonds each period and funds its interest expenditures with lump sum taxes. This could be (for example) because the government wants to supply some bonds to ensure that households have access to a risk-free asset if they want it.

¹³In principle ω could vary over time. I abstract from that here, assuming that in some previous period idiosyncratic shocks divided the households, but that there are no further idiosyncratic shocks.

Let B_t be the real supply of bonds issued in period t . The real government budget constraint is:

$$\frac{B_{t-1}}{\Pi_t} = \frac{B_t}{1+i_t} + (\omega T_t^w + (1-\omega)T_t^p) \quad (10)$$

The government is subject to the interest rate set by the central bank, which is that offered by the good savings product, not the average interest rate received by households. This is because good financial institutions have no transaction costs, buy bonds from the government and then sell to individuals at marginal cost.

The lump sum taxes on the two types of households are always equal, except in the event of the redistributive tax shocks considered in Appendix C.

Bond market clearing implies:

$$B_t = \omega b_t^w + (1-\omega)b_t^p \quad (11)$$

2.4 Closing the model

The firm problem gives an aggregate production function and a Phillips curve (derived in Appendix A.1), and the interest rate on government debt is set by a monetary authority according to a Taylor rule:

$$Y_t = A_t(\omega n_t^w + (1-\omega)n_t^p)^{1-\alpha} \quad (12)$$

$$Y_t(1-\epsilon) + \frac{\epsilon}{1-\alpha} w_t \left(\frac{Y_t}{A_t} \right)^{\frac{1}{1-\alpha}} - \psi(\Pi_t - 1)\Pi_t Y_t + \psi\beta \mathbb{E}_t(\Pi_{t+1} - 1)\Pi_{t+1} Y_{t+1} \left(\frac{C_t}{C_{t+1}} \right)^\sigma = 0 \quad (13)$$

$$1 + i_t = \bar{i} + \left(\frac{\Pi_t}{\bar{\Pi}} \right)^{\delta_\pi} \left(\frac{Y_t}{\bar{Y}} \right)^{\delta_Y} \quad (14)$$

Where A_t is TFP, α is the degree of diminishing returns to scale in the production function, ϵ is the elasticity of substitution between intermediate goods varieties, ψ is the Rotemberg (1982) price adjustment cost parameter, C_t is aggregate consumption, and \bar{i} is the nominal interest rate target.

All output Y_t must be consumed in period t by households, be used to pay price adjustment costs in firms, or be lost to transaction costs in bad financial institutions. The market clearing condition is:

$$Y_t = C_t + \frac{\psi}{2}(\Pi_t - 1)^2 Y_t + \tau \left(\omega(1 - q_t^w)b_t^w + (1-\omega)(1 - q_t^p)b_t^p \right) \quad (15)$$

This completes the model.

2.5 Steady states

There are two steady states of this model. As well as a representative-agent steady state, there is also a steady state with households separated into two types, wealthy (w) households on their Euler equation and poor (p) agents up against their borrowing constraint. This is the same form as the two-agent models in Bilbiie (2008) and others, even though (unlike those existing models) the households have identical preferences and access to asset markets. Poor households do not converge to the positive wealth level of the rich households because their lack of wealth means they do not have an incentive to pay attention to choosing good savings products, and so they face low interest rates, which encourages them not to save.

To see this, consider the model with no costs of attention. All households earn the same interest rate on any savings, and in steady state the Euler equations for the wealthy and poor households respectively are:

$$\frac{1}{1 + \bar{i}} = \beta \quad (16)$$

$$\frac{1}{1 + \bar{i}} = \beta + \bar{\lambda}^p \bar{c}^{p\sigma} \quad (17)$$

From here, \bar{x} denotes the steady state of each variable x_t .

Equations (16) and (17) can only both be true if $\bar{\lambda}^p = 0$, which implies that the poor household is not against the borrowing constraint, and therefore that the households are identical. Existing two agent New Keynesian models obtain steady state inequality either by assuming that the Euler equation doesn't hold for poor households because they are barred from asset markets (e.g. Bilbiie, 2008) or by assuming that the two households have different β parameters (e.g. Bilbiie et al., 2013).

With costly attention, the steady state Euler equations become:

$$\frac{1}{1 + \bar{i}^w} = \frac{\bar{q}^w}{1 + \bar{i}} + \frac{1 - \bar{q}^w}{1 + \bar{i}^b} = \beta \quad (18)$$

$$\frac{1}{1 + \bar{i}^p} = \frac{\bar{q}^p}{1 + \bar{i}} + \frac{1 - \bar{q}^p}{1 + \bar{i}^b} = \beta + \bar{\lambda}^p \bar{c}^{p\sigma} \quad (19)$$

Substituting out for the bad interest rate in terms of the good using equation (1) and combining we obtain:

$$\tau(\bar{q}^w - \bar{q}^p) = \bar{\lambda}^p \bar{c}^{p\sigma} \quad (20)$$

If wealthy households pay more attention to choosing between bond dealers than poor households (i.e. if $\bar{q}^w > \bar{q}^p$) then $\bar{\lambda}^p > 0$ in steady state, which implies that the poor household is at their borrowing constraint. In steady state, the first order condition on

attention is:

$$\bar{b}^h \bar{c}^{h-\sigma} = \mu(1 + \bar{i}^h)^2 \mathcal{I}'(\bar{i}^h), \quad h \in \{w, p\} \quad (21)$$

A poor household at the borrowing constraint has zero assets ($\bar{b}^p = 0$), which in equation (21) implies $\mathcal{I}'(\bar{i}^p) = 0$. This occurs when they pay no attention to their choice of bond dealer, so $\bar{q}^p = g$. As household wealth \bar{b}^h rises the left hand side of equation (21) rises, and so therefore does attention and the probability of choosing a good product \bar{q} .¹⁴ This is because the costs of information do not depend on wealth, but the benefits of being more informed are higher interest rates on the entire stock of wealth saved by the household. If more is saved, the benefits of information are larger (as in Arrow, 1987).

There is therefore a steady state in which wealthy households pay more attention to choosing between bond dealers than poor households who are against their borrowing constraint. Poor households have no ‘skin in the game’, and so pay no attention. This means that they face lower interest rates on average, which reduces their incentive to save. This is the ‘attention trap’, and it gives a two-agent steady state without heterogeneity in household preferences or access to asset markets.

A further implication of equation (20) is that if wealthy households pay less attention to their savings product choices (for example, if the cost μ rises), or if τ falls, then $\bar{\lambda}^p \bar{c}^{p\sigma}$ must fall. $\bar{\lambda}^p$ is the marginal utility of a loosening of the borrowing constraint to a poor household, so a fall in this implies that poor households must be closer to lifting off that constraint, which will be important in Section 3.3. As long as wealthy households process some positive amount of information (as long as μ remains finite and consumption is positive), and there remains some interest rate heterogeneity among bond dealers ($\tau > 0$), then the poor households will remain at the borrowing constraint in steady state ($\bar{\lambda}^p > 0$).

Equation (20) also highlights that the two-agent steady state does not rely on the borrowing constraint being precisely at zero, or on there being a positive supply of government bonds. The key requirement is that the probability of choosing low-price bonds is higher for wealthy households than for poor households against their borrowing constraint. If the model had a zero net supply of bonds and a negative minimum asset limit there would still be an unequal steady state. Poor households would borrow by selling bonds to bond dealers, and wealthy households would save by buying them. Borrowers would process information to try and find the dealers offering the highest price (the lowest interest rate), while wealthy households would aim for a low bond price as in the model above. Just as in this model, poor households would therefore face lower interest rates than wealthy households, which would sustain the wealth inequality.

¹⁴Specifically, attention rises when wealth rises in marginal utility units. In all the exercises considered here that will correspond to increases in b .

The effects of policy would also be qualitatively the same with these alternative assumptions. As poor households lift off their borrowing constraint they have less skin in the borrowing game, so they pay less attention to finding the highest possible bond price, so the interest rate they face rises, just as it does in the model with a positive supply of government debt and no borrowing.

3 Fiscal policy

In this section I show that large fiscal expansions have smaller and more persistent effects on aggregate consumption in the model developed in Section 2 than in a model where inequality is due to heterogeneous discount factors. In both models, large expansions cause poor households to temporarily lift off the borrowing constraint, but in the attention trap model those poor households also then pay more attention to their choice of savings product. They therefore earn higher interest rates, and so they save even more of the windfall from the policy, reducing the impact of that policy on aggregate consumption. This also prolongs the period over which the poor households spend the windfall and return to the borrowing constraint, which makes the effects of the policy more persistent.

If we instead consider small fiscal expansions, which leave poor households at the borrowing constraint every period (as in Bilbiie et al., 2013), then the only difference between the attention trap model and the heterogeneous discount factor model comes from the way the attention of wealthy households reacts to policy. As this is less clear than the attention response of poor households I abstract from this issue by considering the case where information costs are sufficiently low that wealthy households pay enough attention that the probability of choosing a bad product $1 - q^w$ is negligible in all periods.¹⁵ For small fiscal expansions the attention trap and heterogeneous β models are therefore identical, but this is not the case for larger fiscal expansions that cause poor households to temporarily save some of their income.

The specific fiscal expansion studied here is an unanticipated one-period reduction in the lump sum taxes faced by each household, funded by an increase in government debt.¹⁶ The level of real debt decays back to steady state following an AR(1) process

¹⁵For wealthy households, small changes in wealth after a fiscal policy shock are outweighed in the attention decision by changes in the interest rate distribution for most calibrations of the model. In this simplified model, the dispersion of interest rates rises mechanically with the policy rate. However, in Macaulay (2020) I show that this is not how interest rate dispersion behaves in the data. I therefore leave discussion of that result, and cyclical patterns in the attention of wealthy households, to that paper. There I also discuss the interaction of attention with monetary policy shocks, which differs from the case of fiscal policy because of the response of interest rate dispersion.

¹⁶This is one of the policies studied in Bilbiie et al. (2013). I show in Appendix C that the qualitative differences between the attention trap and heterogeneous β model are the same after the other expansion they consider, an intratemporal redistribution from wealthy to poor households.

with persistence ρ .

3.1 The comparison model

To assess the importance of the source of inequality for the effects of fiscal policy I compare the TANK model from Section 2 with a model in which there is no attention problem, and steady state wealth inequality is driven by discount rate heterogeneity (as in Bilbiie et al., 2013). I choose this, rather than the model of limited asset market participation in Bilbiie (2008), to allow the poor households in the comparison model to save positive amounts if the tax windfall is sufficiently large, as they do in the attention trap model.

The model is identical to that in Section 2, except that the cost of information is $\mu = 0$, and a proportion $1 - \omega$ of households have a lower discount factor than other households, $\beta^p < \beta$. The Euler equations of the two household types in steady state together imply:

$$\beta = \beta^p + \bar{\lambda}^p \bar{c}^{\sigma} \quad (22)$$

At any given interest rate, poor households save less than wealthy (high β) households. The only steady state therefore has a two-agent structure similar to that in the model in Section 2, with poor households at the borrowing constraint ($\bar{\lambda}^p > 0$). The key difference is that here wealth inequality is driven by preferences, which are exogenous to fiscal policy, whereas in the attention trap model inequality is maintained by endogenous attention choices. This means that the identity of those at the borrowing constraint is random in the attention trap model, depending only on initial endowments, while here the poor households are poor because they are less patient than others.

The full model equations are given in Appendix A.2. For the comparisons I calibrate this model and the attention trap model using the same values for all common parameters, which means they have the same steady state. The only non-common parameters are β^p and the parameters of the attention problem. I set these such that the discounted interest rate facing poor households in steady state is equal across models ($\beta^p(1 + \bar{i}) = \beta(1 + \bar{i}^p)$), since it is heterogeneity in these discounted rates that generates the steady state inequality in both models. Calibrating the models in this way means that this effect is equally strong in steady state in each model.

3.2 Debt-funded tax cuts

Tax cuts to all households funded by increases in government debt offer a clear view of the differences between the attention trap and heterogeneous β models, as well as being an important part of the way the US government responded to both the Great Recession

and the Covid-19 pandemic.¹⁷ Such tax cuts are expansionary in both models, but if they are sufficiently large that poor households begin to save then the aggregate consumption response is smaller on impact and more persistent in the attention trap model.

3.2.1 Analytic results on MPCs

The differences between the models are driven by the MPC of poor households. I therefore proceed by showing that this MPC is lower in the attention trap model than the heterogeneous β model after large tax cuts. Specifically, I find the partial equilibrium response of poor household consumption to an unanticipated tax cut in period s , holding all aggregate variables constant.¹⁸ Since the two models have the same steady state, I assume that aggregate variables are the same in the two models for this exercise. Household-level variables will differ across models, and the subscript $x \in \{A, B\}$ indicates the attention trap and heterogeneous β models respectively.

Let $\mathcal{B}_{x,s}^p$ be the amount poor households spend on bonds in period s , so $\mathcal{B}_{A,s}^p = \frac{b_{A,s}^p}{1+i_s^p}$ and $\mathcal{B}_{B,s}^p = \frac{b_{B,s}^p}{1+i_s}$. Differentiating the poor household period s budget constraint with respect to taxes gives:

$$\frac{\partial c_{x,s}^p}{\partial T_s^p} = -\frac{1}{1 + \frac{\sigma}{\varphi} w_s^{\frac{1+\varphi}{\varphi}} c_{x,s}^{p-\frac{\sigma+\varphi}{\varphi}}} \left(1 + \frac{\partial \mathcal{B}_{x,s}^p}{\partial T_s^p} \right) \quad (23)$$

First, consider the case of a tax cut that is too small to lift poor households off the borrowing constraint, so $\frac{\partial \mathcal{B}_{x,s}^p}{\partial T_s^p} = 0$ in both models. The derivative $\frac{\partial c_{x,s}^p}{\partial T_s^p}$ is negative, which implies that the MPC of poor households is positive (a tax cut is a fall in T_s^p). It is also the same across both models.

Small tax cuts are therefore expansionary in these models. If all households were on their Euler equation, the MPC out of tax cuts would be zero as Ricardian equivalence would hold, but borrowing constraints mean that poor households are not Ricardian. Their consumption rises substantially, and so therefore does aggregate consumption. Two

¹⁷A lump sum tax cut in the model is a close approximation to the mailing of checks carried out in both circumstances, though the real policies did taper at very high levels of income. While the model exercise is an expansionary policy shock at the steady state, the real fiscal expansions were intended to mitigate large recessions. I abstract from the differential effects of fiscal policy across the business cycle for simplicity, but note that the mechanism differentiating the attention trap and heterogeneous β models is still potentially important in a deep recession. The stimulus packages mentioned were not targeted exclusively at those at borrowing constraints before the recessions began: some of those receiving the stimulus money had only just been pushed to their constraints by the recession, so it is plausible that the stimulus relaxed their constraints. Indeed, Coibion et al. (2020) find that most households intended to mainly use their stimulus checks from the 2020 CARES Act to save or pay down debt.

¹⁸This consumption response is not strictly the marginal propensity to consume once I move to analysing large changes, rather than the gradient of consumption. I maintain the MPC language throughout, however, as it shows the link from this discussion to the wider literature on MPC heterogeneity.

general equilibrium effects amplify the expansion. First, higher consumption means that labour demand rises, and so therefore do wages, further increasing the income and consumption of poor households. Second, in future periods taxes will be high to pay off the increased government debt. Since poor households have a high MPC, their consumption will be below steady state in these periods, which reduces inflation through the Phillips curve. The monetary authority will respond to this by reducing nominal interest rates such that real interest rates are low in those future periods. Anticipating this, when wealthy households observe the tax cut in period s they expect lower future real interest rates, which supports their consumption in period s (see Bilbiie et al. (2013) for a detailed discussion of these channels).

Now turn to the case where the tax cut does lift poor households off their borrowing constraint in period s , so $\frac{\partial \mathcal{B}_{x,s}^p}{\partial T_s^p} < 0$. In both models, this reduces the MPC out of the tax cut because poor households are no longer consuming all of the extra income from the tax cut.¹⁹ The models are not, however, identical here as poor households will increase saving by different amounts.

To see this, consider the simple example where poor households return to the borrowing constraint one period after the tax cut in both models, so $b_{x,s+1}^p = 0$. The Euler equations between periods s and $s + 1$ in the two models are:

$$c_{A,s}^{p-\sigma} = \frac{\beta(1 + \tilde{i}_s^p)}{\Pi_{s+1}} c_{A,s+1}^{p-\sigma}, \quad c_{B,s}^{p-\sigma} = \frac{\beta^p(1 + i_s)}{\Pi_{s+1}} c_{B,s+1}^{p-\sigma} \quad (24)$$

Recall that to compare the models, β^p and the parameters of the information problem are set so that in steady state $\beta(1 + \tilde{i}^p) = \beta^p(1 + \bar{i})$. All aggregate variables are the same in the two models and the tax cut hits at the steady state, so if poor households in the attention trap model did not change their attention we would also have $\beta(1 + \tilde{i}_s^p) = \beta^p(1 + i_s)$.

This will not, however, be the case. We know from equation (9) that when saving rises so does attention, which implies that $\beta(1 + \tilde{i}_s^p) > \beta^p(1 + i_s)$. From this, we have that the consumption response of poor households is lower in the attention trap model.²⁰

The intuition is straightforward. Before the tax cut hits, poor households in the two models have the same incentives to save. When the tax cut comes, they both start to save,

¹⁹Positive shocks therefore imply smaller impact MPCs than negative shocks in both of these models, consistent with the evidence in Bracha and Cooper (2014), Bunn et al. (2018), and Christelis et al. (2019). The smaller impact of fiscal expansions in the attention trap model increases this asymmetry, as the two models are identical for contractionary shocks, just as for small positive shocks.

²⁰The proof is by contradiction. Suppose that $c_{A,s}^p \geq c_{B,s}^p$. From the period s budget constraints, this implies that $\mathcal{B}_{A,s}^p \leq \mathcal{B}_{B,s}^p$. Since $\tilde{i}_s^p < i_s$ by definition, this implies that $b_{A,s}^p < b_{B,s}^p$. In period $s + 1$, poor households therefore have lower asset income in the attention trap model, and so $c_{A,s+1}^p < c_{B,s+1}^p$. Using this and the fact that $\beta(1 + \tilde{i}_s^p) > \beta^p(1 + i_s)$ in the Euler equations (24) implies that $c_{A,s}^p < c_{B,s}^p$, contradicting the initial assumption. The only solution therefore has $c_{A,s}^p < c_{B,s}^p$.

which induces the poor household in the attention trap model to pay more attention to their choice of savings product. In doing so, the interest rate they face rises relative to the policy rate, increasing their desire to save further. They therefore save more, and consume less, in the period of the tax cut than the poor households in the heterogeneous β model. This lower MPC in the attention trap model implies that the fiscal expansion is less powerful in affecting aggregate consumption and output, partly because the direct effect on consumption is smaller, and partly because the general equilibrium effects discussed above are also correspondingly weaker.

The tax cut will have more persistent effects in the attention trap model as long as poor households buy more bonds there (if $b_{A,s}^p > b_{B,s}^p$). In that case, poor household income is higher in period $s + 1$ in the attention trap model due to the extra asset income, and so their consumption is higher than that of the poor households in the heterogeneous β model ($c_{A,s+1}^p > c_{B,s+1}^p$). The intertemporal-MPC of Auclert et al. (2018) is therefore higher in the attention trap model.²¹ If they are consuming more in period $s + 1$, this implies a lower λ_{s+1}^p in equation (7): poor households are closer to staying above their borrowing constraint. In this way larger tax cuts can imply that poor households remain above the borrowing constraint for longer in the attention trap model than in the heterogeneous β model. Intuitively, when poor households save more because they have increased attention, they have more asset income in the next period, which increases their desire to save in that future period. If they do maintain positive savings for multiple periods, they will also maintain positive attention in those future periods, which further supports their saving.

This is not, however, a guaranteed result. The argument above implies that consumption will be lower in the attention trap model in the period of the tax cut, which in turn implies that the amount spent on bonds will be higher ($\mathcal{B}_{A,s}^p > \mathcal{B}_{B,s}^p$). However, since bonds are more expensive for poor households in the attention trap model this is not sufficient to guarantee that the quantity of bonds bought will be higher. For that, we require that the difference in saving budgets $\mathcal{B}_{x,s}^p$ is large relative to the difference in bond prices. Using the definitions of $\mathcal{B}_{x,s}^p$ and \tilde{i}_s^p it can be shown that $b_{A,s}^p > b_{B,s}^p$ if:

$$\mathcal{B}_{A,s}^p > \mathcal{B}_{B,s}^p (1 + \tau(1 - q_s^p)(1 + i_s)) \quad (25)$$

The probability an individual in a poor household chooses a good savings product (q_s^p) is increasing in the amount saved through the first order condition on attention (9). Condition 25 is therefore more easily satisfied for larger tax cuts, which imply larger poor

²¹More generally, this shows that TANK models where poor households leave the borrowing constraint after large positive shocks can predict the positive intertemporal-MPCs found empirically by Auclert et al. (2018), but only for positive income shocks.

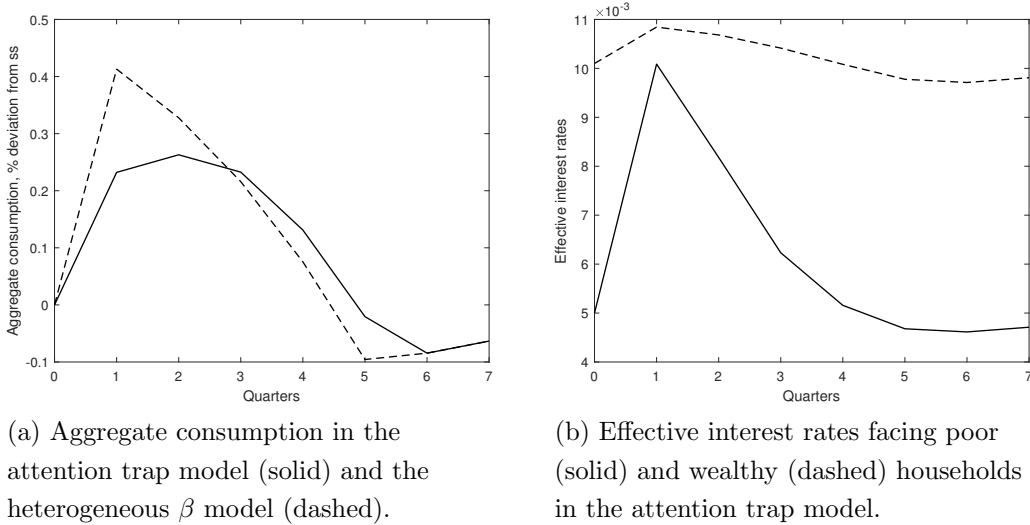
household saving. Even for very large interest rate dispersion (large τ), a sufficiently large tax cut will ensure that $b_{A,s}^p > b_{B,s}^p$, as q_s^p approaches 1 as $b_{A,s}^p$ gets large.

3.2.2 Simulation

The MPC of poor households in the attention trap model is therefore smaller, and for large enough shocks the intertemporal-MPC is larger, than in the heterogeneous β model. Here I simulate a debt-funded aggregate tax cut in the two models, to show the combination of these direct effects and the general equilibrium effects discussed above, in the more general case where poor households continue to save for several periods.

The results are in Figure 1. The shock increases government debt in quarter 1, which then returns to steady state following an AR(1) process with persistence 0.75. This implies a lump sum tax cut in quarter 1, then taxes above steady state from quarter 2 onwards. The initial shock is chosen to be approximately one third of the size of the 2009 fiscal stimulus measures in the US. Details of the calibration are given in Appendix A.3.

Figure 1: Effects of a debt-funded aggregate tax cut



The response of aggregate consumption (Figure 1a) is substantially lower on impact in the attention trap model (solid line) than in the heterogeneous β model (dashed line). With this calibration the rise in aggregate consumption on impact is 44% lower in the attention trap model. Aggregate consumption in the attention trap model is also more persistent, which means that the cumulative aggregate consumption deviation over quarters 1-7 (when the models coincide once more) is only 12% lower in the attention trap model. Poor household saving returns to 0 a quarter later in the attention trap model.

The mechanism can be seen in Figure 1b, which shows the effective interest rates earned by wealthy (dashed line) and poor (solid line) households in the attention trap model. As poor households lift off their borrowing constraint they increase attention

to choosing between savings products. That causes their effective interest rate to rise substantially relative to the rate experienced by the wealthy households, which is the policy rate. The gap between the effective interest rates of the two households widens as debt is brought back to steady state and poor households reduce their saving back to the borrowing constraint, causing them to reduce their attention. This would happen even without further tax changes through the logic of the attention trap, but it is accelerated here by the higher taxes used to repay the government debt, which reduce the income, and so saving, of all households.

While the tax cut causes interest rate heterogeneity to fall in the attention trap model, it is not eliminated entirely. The first order condition on attention (9) implies that the two households only face the same effective interest rate if they save the same amount (in marginal utility units), so no aggregate tax cut of any size will reduce interest rate heterogeneity to zero. In fact, the aggregate tax cut in this simulation the heterogeneity in saving in marginal utility units, because wealthy households absorb the bulk of the extra government debt, and they do so by reducing consumption.

3.3 Factors affecting the power of fiscal policy

As well as implying that fiscal expansions have smaller and more persistent effects, the attention trap model differs from the heterogeneous β model in that a number of features of the economic environment affect the power of fiscal policy through novel channels not present in that comparison model. In general, anything that affects the incentives of poor households to pay attention to their choice between savings products will affect their consumption response to tax cuts, and so the aggregate impact of such a policy shock. I analyse three such factors: the cost of information, asset income taxation, and interest rate dispersion.

3.3.1 The cost of information

The cost of information μ only enters the equilibrium conditions of the model in the attention first order condition (9). Using the definitions of i^b (1) and \mathcal{I} (3), this can be written as:

$$b_t^h c_t^{h-\sigma} = \frac{\mu}{\tau} \log \left(\frac{q_t^h (1-g)}{(1-q_t^h)g} \right) \quad (26)$$

Taking the same approach as in Section 3.2.1, consider the response of the probability of choosing a good savings product to a change in taxes in period s , holding all aggregate

variables constant:

$$\frac{\partial q_s^h}{\partial T_s^h} = \frac{c_s^{h-\sigma} q_s^h (1 - q_s^h) \tau}{\mu} \left(\frac{\partial b_s^h}{\partial T_s^h} - \frac{\sigma b_s^h}{c_s^h} \frac{\partial c_s^h}{\partial T_s^h} \right) \quad (27)$$

For wealthy households $q_s^w \approx 1$, so $\frac{\partial q_s^w}{\partial T_s^w} \approx 0$. For poor households, however, equation (27) shows that the elasticity of q_s^p to tax cuts is higher when the cost of attention μ is low.

Intuitively, poor households pay no attention in steady state if the cost is positive, as they derive no benefit from information while they are not saving. When they start to save, they increase attention until the (now positive) marginal benefits of doing so equal the marginal costs. With a lower marginal cost of information, poor households therefore increase their attention by more after a tax cut. Their effective interest rate rises by more, and so they save more, which further reduces their MPC out of the tax cut and increases the intertemporal-MPC, making the consumption response to the tax cut even smaller and more persistent.²² These effects are amplified by the fact that when poor households save more and consume less in the period of the shocks, that appears in equation (27) as a more negative $\frac{\partial b_s^p}{\partial T_s^p}$ and a less negative $\frac{\partial c_s^p}{\partial T_s^p}$: the extra saving adds to the incentives to increase attention.

3.3.2 Asset income taxes

Asset income taxation can be incorporated into the attention trap model by rewriting the budget constraint as:²³

$$c_t^h + \frac{b_t^h}{1 + \tilde{i}_t^h} = \frac{b_{t-1}^h}{\Pi_t} (1 - T_b) + w_t n_t^h + \mathcal{D}_t - T_t^h \quad (28)$$

The asset income tax T_b does not directly alter the attention first order condition, but the Euler equation becomes:

$$c_t^{h-\sigma} = \beta (1 + \tilde{i}_t^h) (1 - T_b) \mathbb{E}_t \left(\frac{c_{t+1}^{h-\sigma}}{\Pi_{t+1}} \right) + (1 + \tilde{i}_t^h) \lambda_t^h \quad (29)$$

²²This, of course, is only one effect of lower information costs. Lower μ also implies that when poor households do save, less output is wasted in the transaction costs of bad bond dealers as poor households choose them with a lower probability. I study the implications of policies that lower information costs in a more general setting with profit maximising banks and cyclical variation in the attention of all households in Macaulay (2020).

²³The revenue from the asset income tax also enters the government budget constraint. A debt-funded tax cut is therefore cheaper for the government because with more asset supply they increase the revenue from this tax. This does affect the impacts of a debt-funded lump sum tax cut, but here I focus on the channels that are novel to the attention trap model.

Any rise in the effective interest rate \tilde{i}_t^h therefore has a smaller effect on the intertemporal path of consumption when asset income taxes are higher. This means that after a debt-funded lump sum tax cut, a given rise in poor household attention will lead to less of a rise in poor household saving. With less saving, poor households increase their attention by less, so the asset tax indirectly affects the attention first order condition. Through both of these channels, asset income taxation therefore implies that lump sum debt-funded tax cuts will have larger and less persistent effects.

This suggests that cuts to asset income taxes, as occurred in the UK in 2016, could reduce the power of fiscal policy to stimulate consumption. This analysis, however, abstracts from the effect of asset income taxation on wealthy households, in part because I assume $q^w \approx 1$, and also because the tax is levied on the payoff from bonds while attention acts to lower their price. An alternative setup would have all bonds sold at the same price, with attention increasing the interest rate paid out in the next period. In that case, lower taxes on interest income directly increase the returns to attention, which leads to an additional effect of lower T_b . If wealthy households are not already paying so much attention that $q^w \approx 1$, then lower T_b leads them to pay more attention in every period, which reduces the output wasted on transaction costs at bad bond dealers.²⁴ In this way reducing asset income taxes may be beneficial for welfare, though that effect is absent in the model used here.

3.3.3 Interest rate dispersion

Interest rate dispersion is controlled in the model by the parameter τ . The direct effects of dispersion on attention can be seen in equations (26) and (27): lower dispersion (lower τ) implies a lower level of attention and a smaller response of attention to taxes. This is because lower interest rate dispersion reduces the interest rate gain from choosing a good savings product rather than a bad, and so reduces the utility benefit of attention. Just as lower marginal costs of attention increase the level and response of attention to tax cuts for poor households, the lower marginal benefit of attention brought about by lower dispersion reduces both q_t^p and $\frac{\partial q_s^p}{\partial T_s^p}$, if everything else is held constant.

This is not, however, the only effect of a fall in interest rate dispersion. Recall that from equation (20), the steady state shadow value of relaxing the borrowing constraint for poor households ($\bar{\lambda}^p$) falls with τ . The savings poor households would have chosen without that constraint rise closer to zero.

To see why, notice that since wealthy households pay full attention ($q^w \approx 1$) the steady state policy rate, determined by their Euler equation, is unaffected by τ : $1 + \bar{i} = \beta^{-1}$. Lower dispersion must translate in steady state to a higher interest rate on the bad

²⁴These alternative assumptions do not change the qualitative results in the rest of this paper.

products \bar{i}^b . Lower dispersion therefore leads to a higher effective interest rate for any level of attention that implies bad products are chosen with positive probability. Poor households therefore face higher effective interest rates in steady state, and this increases their incentives to save.

As discussed in Section 2.5, this does not lead to poor households saving in steady state unless τ reaches 0. An intermediate fall in τ does however imply that poor households begin to save after smaller tax cuts. Again consider the partial equilibrium effects of a tax cut in period s : substituting the budget constraint and the definition of \tilde{i}_s^p into the Euler equation of a poor household beginning to save gives:

$$T_s^p + \mathcal{B}_s^p = w_s n_s^p + \mathcal{D}_s - \left(\frac{\Pi_{s+1}}{\beta(1+i_s)} + \frac{\tau(1-q_s^p)\Pi_{s+1}}{\beta} \right)^{\frac{1}{\sigma}} c_{s+1}^p \quad (30)$$

To find the maximum level of taxes such that the no-borrowing constraint does not bind for poor households, take this and set $\mathcal{B}_s^p = 0$, which implies $q_s^p = g$. A fall in τ causes this maximum tax to rise, implying that the tax cut required to lift poor households from their borrowing constraint is smaller. If a tax cut is just small enough to leave households at the borrowing constraint with some high level of τ , a fall in dispersion will therefore reduce the power of such a tax cut as the poor household will start to save.

For a given tax cut that is large enough to cause poor households to start saving even with a high τ , however, the effects of lower dispersion are more complicated. In equation (30), a lower τ directly increases the amount saved for a given T_s^p . However, lower τ also implies less attention for any given level of saving through equation (26), reducing q_s^p , which in equation (30) acts to reduce saving.

With the calibration used in Section 3.2.2, reducing τ increases $\tau(1-q_s^p)$. The effective interest rate faced by poor households in the period of the tax cut (\tilde{i}_s^p) is therefore lower when τ is lower, even though for any given level of attention the effective interest rate would be higher. However, this does not lead to a greater MPC when interest rate dispersion is lower, because period s consumption is determined by the discounted sum of future effective interest rates until the poor household returns to the borrowing constraint, not just the effective interest rate in the period of the tax cut.


As taxes rise to pay off the government debt issued to fund the tax cut, poor households reduce their saving, and so reduce their attention to choosing between savings products. This fall in attention has a smaller effect on the effective interest rate when interest rate dispersion is low, and so in some period after the tax cut the effective interest rate in the low-dispersion environment overtakes that in a high-dispersion setting. The discounted sum of effective interest rates can therefore be higher when dispersion is low even if the initial interest rate is lower, reducing the MPC of poor households and the

impact of the tax cut.

Moreover, low interest rate dispersion implies that poor households remain above the borrowing constraint for longer, because the minimum income required for them to save is lower (equation (30)). Positive saving in some future period reduces expected consumption in that period, which also causes poor households to reduce initial consumption and increase saving in order to smooth consumption.

These effects are shown in Figure 2, which repeats the baseline simulation of Figure 1 (solid line) and compares it to a simulation with a 50% lower τ (dashed line). With this calibration, the effective interest rate faced by poor households is higher in the low- τ case in all periods except period 1, when the tax cut hits. Poor households therefore increase their consumption by less when they receive the tax cut in the low- τ case, so the tax cut has an 18% smaller effect on aggregate consumption on impact than in the baseline with higher interest rate dispersion.²⁵

Figure 2: Effects of a debt-funded aggregate tax cut, high vs low τ



(a) Aggregate consumption in the attention trap model with high (solid) and low (dashed) τ .

(b) Effective interest rates facing poor households with high (solid) and low (dashed) τ .

These findings are of particular relevance in the current environment of persistently low natural rates of interest (Holston et al., 2017), since low nominal interest rates are likely to coincide with low interest rate dispersion, as banks are unwilling to pass on negative interest rates to households (Eggertsson et al., 2019).²⁶ Low interest rate dispersion may therefore reduce the power of fiscal policy precisely when it is needed to

²⁵The result that lower τ reduces the impact of fiscal policy is robust to various different calibrations. Specifically, as long as government debt decays with a persistence $\rho > 0.05$, b^p remains above 0 for at least one more quarter in the low- τ environment, so c^p responds less to the tax cut when τ is low.

²⁶While Eggertsson et al. (2019) study interest rates in Sweden, this effect is not confined to economies with negative policy rates. For the UK, Barrett (2018) notes in the Financial Times that “as interest rates continue to bump along the bottom, the difference between the best and the worst easy-access savings rates are not that great”.

play a greater role in stabilising demand in recessions, that is when monetary policy is constrained by the effective lower bound on policy rates.

4 Evidence for attention rising after a tax cut

Fiscal policy has smaller and more persistent effects in the attention trap model because poor households increase their attention to choosing between savings products when their income rises after a tax cut. In this section I provide evidence for such an attention response, using cross-state variation to identify the effect of the Tax Cuts and Jobs Act (TCJA) in the US on a survey-based measure of uncertainty about savings interest rates. The TCJA was signed on 22nd December 2017 and came into effect on 1st January 2018, and the largest single tax change was the cut in the federal corporate income tax rate from 35% to 21%.²⁷ I find that this significantly reduced uncertainty about savings among poorer households, consistent with the model prediction that these households increased their attention.

4.1 Data and identification

The key data source is the Survey of Consumer Expectations (SCE). This large panel survey asks respondents for the probability that the interest rate on savings will rise in the next 12 months. To construct a measure of uncertainty about savings, I note that since “interest rates will rise” is a binary statement, the variance in individual i 's belief about the statement is given by $p_i(1 - p_i)$, where p_i is their subjective probability of the statement being true, and their survey response. I use this variance as a measure of how uncertain individual i is about the savings market. Intuitively, the further an individual's response is from 50/50 the less uncertain they are about the path of future savings rates.

In the model, tax cuts cause poor households to pay more attention to their choice of savings product, reducing their uncertainty about which product to choose for their saving in that period. This is not exactly the same as the uncertainty measured using the SCE, which is uncertainty about the path of future savings interest rates rather than current interest rates. However, the two forms of uncertainty are very likely to be correlated, because in reality households can choose between assets of different maturities. This means that information about the future path of interest rates is valuable for comparing currently available assets. Greater attention to current asset choice will therefore involve

²⁷The corporate tax cut was the largest single tax change in the TCJA, but other taxes changed as well. The 19 responsive states in my treatment group contain 4 of the 5 states most exposed to the income tax changes in the act as estimated by Kumar (2020), using the method of Zidar (2019).


reducing uncertainty about the future path of interest rates.²⁸

To identify the effect of the TCJA on this measure of uncertainty, I make use of the results on the state-level effects of corporate tax cuts in Liu and Williams (2019) (LW). They find that corporate tax cuts only lead to a significant increase in GDP in 19 of the states of the US, while most states see no significant effect. This heterogeneity is explained by differences in the industry make-up of different states. In particular, less capital intensive states are more likely to expand after a federal corporate tax cut.

I split states following LW’s results. I refer to the states that see a significant expansion in GDP after corporate tax cuts as ‘responsive’,²⁹ and the other states as ‘unresponsive’. Since it is income that determines attention in the model, the model predicts that uncertainty will fall by more in states where the tax cut causes an expansion in economic activity, that is in the responsive states. Industry and capital intensity are slow-moving variables, so should not change substantially in the months immediately before and after the tax cut. The LW estimation uses data from 1950-2017, stopping before the TCJA.

Figure 3 plots the average uncertainty in the responsive and unresponsive states over time, with and without controls.³⁰ In both graphs, there is no clear difference in the trends in uncertainty between the two groups before the tax cut. After the tax cut, uncertainty falls temporarily in the responsive states relative to the unresponsive states. The difference closes up after just over a year.

Figure 3: Uncertainty in unresponsive and responsive states. The tax cut occurs in 2018m1.



(a) Raw series

(b) With controls

²⁸See Yankov (2018) for evidence on the demand for certificates of deposit of various maturities. Yankov (2018) also finds that deposit rates are very persistent, which means that processing information about the current distribution of interest rates will automatically reduce uncertainty about the future rate distribution, as in other information processing problems with persistent states of the world (see e.g. Steiner et al., 2017).

²⁹These are CO, FL, GA, HI, ID, IL, MA, MD, ME, MI, MO, NE, NH, NJ, NV, RI, SC, WA, WV.

³⁰Specifically, the series in Figure 3b is the residual uncertainty after controlling for month-of-the-year, state and individual fixed effects, log income, and financial fragility, where financial fragility is defined as the self-reported probability of the household being unable to pay their bills in the next 3 months.

In Section 4.2 I formally test for this fall in uncertainty in responsive states. I find that it is significant, and is concentrated among households with low incomes but who are not in immediate danger of being unable to pay their bills. This is consistent with the model, in which a tax cut causes a temporary rise in attention from poor households who increase their saving a little after the cut. I also test, and fail to reject, the assumption of parallel trends in uncertainty before the tax cut.

4.2 Testing for the effect of the TCJA on uncertainty

Figure 3 suggests that uncertainty about interest on savings fell in responsive states relative to unresponsive states after the tax cut. To test this, I estimate:

$$unc_{ist} = D_{st} \sum_{c=1}^4 \beta_c \mathbb{1}_{C_{ist}=c} + \gamma_i + \delta_s + \zeta_t + \theta X_{ist} + \varepsilon_{ist} \quad (31)$$

The dependent variable unc_{ist} is the uncertainty in the survey response of individual i in state s at time t . D_{st} is the treatment dummy: it equals 1 if the individual lives in a state that is responsive to corporate tax cuts and the date of the survey is on or after the date of the tax cut coming into force (January 2018).³¹

I allow the tax cut to have different effects depending on whether the individual has high or low income (above or below the median in the sample), and if they are financially fragile, which I define as reporting a greater than 5% chance of being unable to pay their bills in the next 3 months. C_{ist} is a categorical variable for the combination of these two characteristics. The model predicts that the households who change their attention after the tax cut are those who are close to their borrowing constraint, so if the model is a good representation of attention decisions in this market uncertainty should fall only for those with low income who are not financially fragile.³²

The remaining terms in equation (31) are individual, state, and time fixed effects, and time-varying individual-level controls (log income and financial fragility). Individual fixed effects account for any demographic variables that do not change over the time the individual remains in the survey, though there is little demographic difference between responsive and unresponsive states (see Appendix D.1). I estimate equation (31) both

³¹I do not include leads and lags of the treatment variable in the regression because individuals only remain in the survey for an average of 7 months, so leads and lags would substantially reduce the number of observations available. I therefore leave the testing of the parallel trend assumption to another specification which does not make as much use of the panel dimension of the survey (see Figure 4).

³²Intuitively, if there is a substantial probability that a household won't be able to pay their bills, then they aren't likely to be close to starting to save, with or without the tax cut. Low income non-fragile households are the ones that could increase saving when the tax cut hits, but who aren't already paying lots of attention. They are therefore the relevant group for the mechanism of the paper.

over the whole survey with heteroskedasticity-robust standard errors, and over the slightly smaller set of survey respondents who do not move state while in the survey, for which I can cluster standard errors at the state level.³³ The results are displayed in Table 1.

Table 1: Regression on the uncertainty in forecasts of the interest rate on savings.

	(1)	(2)
	Uncertainty	Uncertainty
Low income non-fragile	-0.00887** (-2.46)	-0.00854** (-2.36)
High income non-fragile	-0.00145 (-0.36)	-0.000903 (-0.28)
Low income fragile	0.00253 (0.71)	0.00331 (0.86)
High income fragile	-0.00416 (-0.94)	-0.00361 (-0.89)
Individual FE	Yes	Yes
State FE	Yes	No
Time FE	Yes	Yes
Clustered SE	No	Yes
Observations	102205	100241

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

As predicted by the model, low income non-fragile households become less uncertain after the tax cut in responsive states relative to unresponsive states, and there is no effect on the uncertainty of high income or fragile households. The fall in the uncertainty of low income non-fragile households is non-negligible: it is 6.4% of the mean uncertainty for these households across the survey, or 10.0% of the standard deviation.

To test for parallel trends, I first regress savings interest rate uncertainty on individual fixed effects and the time-varying individual characteristics X_{ist} . I then compute the average of the residual uncertainty by state, high/low income, and fragile/non-fragile. This gives four continuous series of average uncertainty for each state, one for each combination of income and financial fragility. Using these average uncertainty series, I estimate:

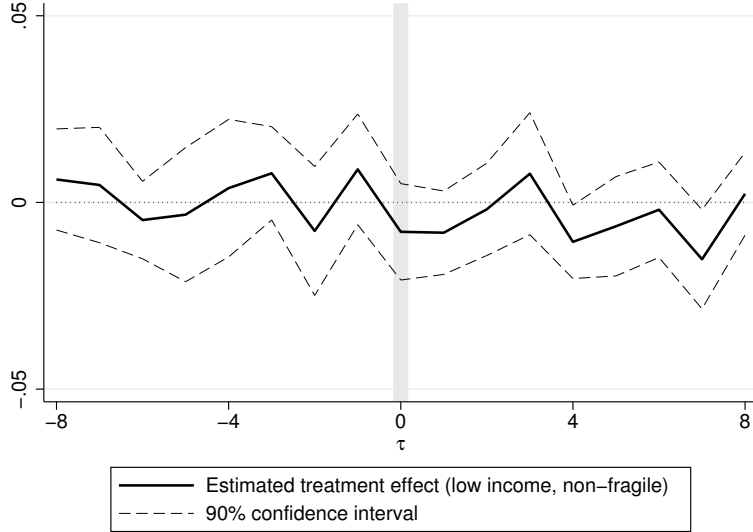
$$u\bar{n}c_{cst} = \sum_{\tau=-8}^8 \sum_{c=1}^4 \tilde{D}_{st-\tau} \beta_{c\tau} \mathbb{1}_{C_{cst-\tau}=c} + \gamma_c + \delta_s + \zeta_t + \varepsilon_{cst} \quad (32)$$

Here $u\bar{n}c_{cst}$ is the average residual uncertainty for individuals with income and fragility

³³Without movers individual fixed effects ensure all state fixed effects are zero.

characteristics c in state s at time t . \tilde{D}_{st} equals 1 in states that are responsive to corporate tax cuts on the date of the cut, and equals zero for all states before and after.³⁴ The index τ therefore denotes leads and lags of the treatment. C_{cst} denotes the combination of income and financial fragility characteristics as before. The remaining terms are characteristic, state and time fixed effects. The results are plotted in Figure 4.³⁵

Figure 4: Estimated coefficients $\beta_{c\tau}$ in equation (32) for $c = (\text{low income, non-fragile})$.



I do not reject the parallel trend assumption. The effect of leads of the tax cut is insignificantly different from zero. After the tax cut, the estimated effects of the policy on uncertainty become negative, as in the individual-level model in equation (31). The magnitudes are similar to the individual-level model, with the effects on low income non-fragile individuals 4 and 7 months after the tax cut estimated at -0.011 and -0.015 respectively. These lags of the policy are significant at the 10% level, despite the effects of the tax cut being estimated much less precisely in this specification than in equation (31), since aggregating to the state-characteristic level reduces the number of observations by more than a factor of 10.

Taken together, this evidence suggests that the tax cut was associated with lower uncertainty about savings markets for individuals with low income but who are not financially fragile, in states that are responsive to corporate tax cuts. This is consistent with the model, and provides evidence for its key mechanism. Further tests in Appendix

³⁴This differs from the treatment dummy in equation (31), D_{st} , which is 1 for all periods on or after the date of the tax cut. This change aids the interpretation of the coefficients on the leads and lags of \tilde{D}_{st} .

³⁵For brevity, Figure 4 presents only the coefficients for the treatment on the group of interest (low income non-fragile). The regression is weighted using the number of observations in the average used to construct each cst cell, and standard errors are clustered at the state level.

D.2 also support the model: the tax cut is associated with an increase in planned saving for low income non-fragile households, but is not associated with a change in the subjective probability of an interest rate rise. The results here could therefore be driven by households saving more, but are not driven by changes in the direction of expected interest rate moves - only the uncertainty around that expectation.

5 Conclusion

The existence of wealth inequality alters many of the predictions of macroeconomic models, with fiscal policy a particularly strong example. I have shown that the source of the inequality also affects those predictions. In a model where wealth inequality is driven by endogenous asset return heterogeneity, fiscal expansions have smaller and more persistent effects on consumption than in a comparable two-agent model in which inequality comes from heterogeneous discount factors. Lower costs of information about asset returns, taxes on asset income, and interest rate dispersion all reduce the impact of fiscal expansions further in this model.

These results are due to a complementarity between saving and information processing. As poor households increase their saving after a fiscal expansion, they have greater incentives to pay attention to their choice of savings product. They make better savings choices, and the interest rate they face rises relative to the policy rate, causing them to save more and consume less. Cross-state variation in uncertainty about savings interest rates after the 2017 Tax Cuts and Jobs Act provides evidence for this response of attention to tax cuts.

References

- Acharya, S. and Dogra, K. (2020). Understanding HANK: Insights From a PRANK. *Econometrica*, 88(3):1113–1158.
- Acharya, S. and Wee, S. L. (2020). Rational inattention in hiring decisions. *American Economic Journal: Macroeconomics*, 12(1):1–40.
- Adam, K. (2009). Monetary policy and aggregate volatility. *Journal of Monetary Economics*, 56(SUPPL.):S1–S18.
- Afrouzi, H. (2020). Strategic Inattention, Inflation Dynamics, and the Non-Neutrality of Money.
- Almgren, M., Gallegos Dago, J. E., Kramer, J., and Lima, R. (2019). Monetary Policy and Liquidity Constraints: Evidence from the Euro Area.

- Arrow, K. J. (1987). The demand for information and the distribution of income. *Probability in the Engineering and Informational Sciences*, 1(1):3–13.
- Auclert, A., Rognlie, M., and Straub, L. (2018). The Intertemporal Keynesian Cross. *NBER Working Paper Series*, 25020.
- Barrett, C. (2018). What would it take for you to switch your bank account? *Financial Times*, 27th July.
- Bayer, C., Born, B., and Luetticke, R. (2020). The Liquidity Channel of Fiscal Policy. *CEPR Discussion Papers*, 14883.
- Benhabib, J., Bisin, A., and Luo, M. (2019). Wealth distribution and social mobility in the US: A quantitative approach. *American Economic Review*, 109(5):1623–1647.
- Best, J. and Dogra, K. (2019). Capital Management and Inequality.
- Bilbiie, F. O. (2008). Limited asset markets participation, monetary policy and (inverted) aggregate demand logic. *Journal of Economic Theory*, 140(1):162–196.
- Bilbiie, F. O. (2019a). Monetary Policy and Heterogeneity: An Analytical Framework.
- Bilbiie, F. O. (2019b). The New Keynesian cross. *Journal of Monetary Economics*.
- Bilbiie, F. O., Monacelli, T., and Perotti, R. (2013). Public Debt and Redistribution with Borrowing Constraints. *The Economic Journal*, 123(566):F64–F98.
- Blanchard, O. and Summers, L. H. (2019). Introduction. In Blanchard, O. and Summers, L. H., editors, *Evolution or Revolution? Rethinking Macroeconomic Policy after the Great Recession*. MIT Press.
- Bracha, A. and Cooper, D. (2014). Asymmetric responses to income changes: The payroll tax increase versus tax refund in 2013. *Economics Letters*, 124(3):534–538.
- Briggs, J., Cesarini, D., Lindqvist, E., and Östling, R. (2020). Windfall Gains and Stock Market Participation. *Journal of Financial Economics*.
- Bunn, P., Le Roux, J., Reinold, K., and Surico, P. (2018). The consumption response to positive and negative income shocks. *Journal of Monetary Economics*, 96:1–15.
- Campanale, C. (2007). Increasing returns to savings and wealth inequality. *Review of Economic Dynamics*, 10(4):646–675.
- Cantore, C. and Freund, L. B. (2021). Workers, capitalists, and the government: fiscal policy and income (re)distribution. *Journal of Monetary Economics*.
- Christelis, D., Georgarakos, D., Jappelli, T., Pistaferri, L., and Van Rooij, M. (2019). Asymmetric consumption effects of transitory income shocks. *Economic Journal*, 129(622):2322–2341.

- Coibion, O., Gorodnichenko, Y., and Weber, M. (2020). How Did U.S. Consumers Use Their Stimulus Payments? *NBER Working Paper Series*, 27693.
- Congressional Budget Office (2009). cost estimate for the conference agreement for H.R. 1, the American Recovery and Reinvestment Act of 2009.
- Dasgupta, K. and Mondria, J. (2018). Inattentive importers. *Journal of International Economics*, 112:150–165.
- Debortoli, D. and Galí, J. (2018). Monetary policy with heterogeneous agents: Insights from TANK models. *Economics Working Papers, Department of Economics and Business, Universitat Pompeu Fabra*, 1686.
- Eggertsson, G. B., Juelsrud, R. E., Summers, L. H., and Wold, E. G. (2019). Negative nominal interest rates and the bank lending channel. *NBER Working Paper Series*, 25416.
- Fagereng, A., Guiso, L., Malacrino, D., and Pistaferri, L. (2020). Heterogeneity and Persistence in Returns to Wealth. *Econometrica*, 88(1):115–170.
- Galí, J., López-Salido, J. D., and Vallés, J. (2007). Understanding the effects of government spending on consumption. *Journal of the European Economic Association*, 5(1):227–270.
- Giambattista, E. and Pennings, S. (2017). When is the government transfer multiplier large? *European Economic Review*, 100:525–543.
- Gornemann, N., Kuester, K., and Nakajima, M. (2016). Doves for the Rich, Hawks for the Poor? Distributional Consequences of Monetary Policy. *CEPR Discussion Papers*, 11233.
- Hagedorn, M., Manovskii, I., and Mitman, K. (2019). The Fiscal Multiplier. *NBER Working Paper Series*, 25571.
- Holm, M. B., Paul, P., and Tischbirek, A. (2020). The Transmission of Monetary Policy under the Microscope. *Federal Reserve Bank of San Francisco Working Paper Series*, 2020-03.
- Holston, K., Laubach, T., and Williams, J. C. (2017). Measuring the natural rate of interest: International trends and determinants. *Journal of International Economics*, 108:S59–S75.
- Kacperczyk, M., Nosal, J., and Stevens, L. (2019). Investor sophistication and capital income inequality. *Journal of Monetary Economics*, 107:18–31.
- Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary policy according to HANK. *American Economic Review*, 108(3):697–743.

- Kaplan, G. and Violante, G. L. (2018). Microeconomic heterogeneity and macroeconomic shocks. *Journal of Economic Perspectives*, 32(3):167–194.
- Kumar, A. (2020). Did Tax Cuts and Jobs Act Create Jobs and Stimulate Growth? Early Evidence Using State-Level Variation in Tax Changes. *Federal Reserve Bank of Dallas Working Paper Series*, 2001.
- Lagarde, C. (2020). Opening remarks at the EUI’s State of the Union event.
- Lei, X. (2019). Information and Inequality. *Journal of Economic Theory*, 184:104937.
- Liu, C. and Williams, N. (2019). State-level implications of federal tax policies. *Journal of Monetary Economics*, 105:74–90.
- Lusardi, A., Michaud, P. C., and Mitchell, O. S. (2017). Optimal financial knowledge and wealth inequality. *Journal of Political Economy*, 125(2):431–477.
- Macaulay, A. (2020). Cyclical Attention to Saving. *UniCredit Foundation Working Papers Series*, 140.
- Maćkowiak, B. and Wiederholt, M. (2009). Optimal Sticky Prices under Rational Inattention. *American Economic Review*, 99(3):769–803.
- Maćkowiak, B. and Wiederholt, M. (2015). Business Cycle Dynamics under Rational Inattention. *Review of Economic Studies*, 82(4):1502–1532.
- Malmendier, U. and Nagel, S. (2011). Depression Babies: Do Macroeconomic Experiences Affect Risk Taking? *The Quarterly Journal of Economics*, 126(1):373–416.
- Matějka, F. and McKay, A. (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–298.
- McKay, A. (2013). Search for financial returns and social security privatization. *Review of Economic Dynamics*, 16(2):253–270.
- McKay, A., Nakamura, E., and Steinsson, J. (2016). The power of forward guidance revisited. *American Economic Review*, 106(10):3133–3158.
- Occhino, F. (2019). The Macroeconomic Effects of the Tax Cuts and Jobs Act. *Federal Reserve Bank of Cleveland Working Paper*, 19-28.
- Powell, J. (2020). FOMC Press Conference, July 29 2020.
- Rotemberg, J. J. (1982). Sticky Prices in the United States. *Journal of Political Economy*, 90(6):1187–1211.
- Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690.

- Sims, C. A. (2006). Rational inattention: Beyond the linear-quadratic case. *American Economic Review*, 96(2):158–163.
- Steiner, J., Stewart, C., and Matějka, F. (2017). Rational Inattention Dynamics: Inertia and Delay in Decision-Making. *Econometrica*, 85(2):521–553.
- Werning, I. (2015). Incomplete Markets and Aggregate Demand. *NBER Working Paper Series*, 21448.
- Yankov, V. (2018). In Search of a Risk-Free Asset: Search Costs and Sticky Deposit Rates.
- Zidar, O. (2019). Tax cuts for whom? Heterogeneous effects of income tax changes on growth and employment. *Journal of Political Economy*, 127(3):1437–1472.

A Model equations

A.1 Taylor rule and firm problem in the attention trap model

A perfectly competitive final goods producer combines varieties of intermediate goods $Y_t(i)$ to make the final good using a standard CES aggregator:

$$Y_t = \left(\int_0^1 Y_t(i)^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}} \quad (33)$$

Cost minimisation gives the demand curve facing an intermediate goods producer i :

$$Y_t(i) = \left(\frac{P_t(i)}{P_t} \right)^{-\epsilon} Y_t \quad (34)$$

The price index P_t is also the final good producer's marginal cost, and so the price offered to consumers:

$$P_t = \left(\int_0^1 P_t(i)^{1-\epsilon} \right)^{\frac{1}{1-\epsilon}} \quad (35)$$

Intermediate goods firms are monopolistically competitive, face Rotemberg (1982) quadratic adjustment costs in prices, and have the production function:

$$Y_t(i) = A_t N_t(i)^{1-\alpha} \quad (36)$$

Here $N_t(i)$ is labour employed by the intermediate firm i . Labour supplied from one household is a perfect substitute for labour from any other household.

The firm chooses their price to maximise the expected discounted sum of profits, subject to this production function and demand:

$$\max_{P_t(i)} \mathbb{E}_0 \sum_{t=0}^{\infty} \Lambda_{0,t} \left[P_t(i) Y_t(i) - W_t N_t(i) - \frac{\psi}{2} \left(\frac{P_t(i)}{P_{t-1}(i)} - 1 \right)^2 P_t Y_t \right] \quad (37)$$

subject to

$$Y_t(i) = \left(\frac{P_t(i)}{P_t} \right)^{-\epsilon} Y_t \quad (38)$$

$$Y_t(i) = A_t N_t(i)^{1-\alpha} \quad (39)$$

Firm ownership is equally distributed among households, who contract out firm operation to an agent with a stochastic discount factor of $\Lambda_{t,t+1} = \frac{\beta}{\bar{\Pi}_{t+1}} \left(\frac{C_t}{C_{t+1}} \right)^\sigma$, where C_t is aggregate consumption.³⁶ All profits are distributed back to households equally.

The first order condition of an intermediate goods producer gives the Phillips curve when we substitute in that all intermediate firms are identical, so make identical choices:

$$Y_t(1-\epsilon) + \frac{\epsilon}{1-\alpha} w_t \left(\frac{Y_t}{A_t} \right)^{\frac{1}{1-\alpha}} - \psi(\Pi_t - 1)\Pi_t Y_t + \psi\beta \mathbb{E}_t(\Pi_{t+1} - 1)\Pi_{t+1} Y_{t+1} \left(\frac{C_t}{C_{t+1}} \right)^\sigma = 0 \quad (40)$$

The production function of intermediate producers gives us an equation for aggregate production, since all firms make identical choices:

$$Y_t = A_t(\omega n_t^w + (1-\omega)n_t^p)^{1-\alpha} \quad (41)$$

Monetary policy is conducted according to a Taylor rule:

$$1 + i_t = \bar{i} + \left(\frac{\Pi_t}{\bar{\Pi}} \right)^{\delta_\Pi} \left(\frac{Y_t}{\bar{Y}} \right)^{\delta_Y} \quad (42)$$

A.2 Comparison model equations

The household side of the model consists of an Euler equation, intratemporal labour supply condition, and a budget constraint for each type of household, $h \in \{w, p\}$:

$$c_t^{h-\sigma} = \beta^h (1 + i_t) \mathbb{E}_t \left(\frac{c_{t+1}^{h-\sigma}}{\Pi_{t+1}} \right) + (1 + i_t) \lambda_t^h \quad (43)$$

³⁶This is similar to the assumption in the HANK model in Debortoli and Galí (2018), where the relevant stochastic discount factor is determined by the average consumption of the households who own equity, weighted by their equity holdings. Here, equity is equally distributed among households and fixed, hence the use of aggregate consumption.

$$c_t^{h\sigma} n_t^{h\varphi} = w_t \quad (44)$$

$$c_t^h + \frac{b_t^h}{1+i_t} = \frac{b_{t-1}^h}{\Pi_t} + w_t n_t^h + \mathcal{D}_t - T_t^h \quad (45)$$

The firm, policy and market clearing blocks of the model are as in Section 2 and Appendix A.1, except that there are no bad bond dealer transaction costs in the goods market clearing condition. The equations consist of a Phillips curve, production function, Taylor rule, government budget constraint, tax rule, and market clearing conditions for bonds and goods:

$$Y_t(1-\epsilon) + \frac{\epsilon}{1-\alpha} w_t \left(\frac{Y_t}{A_t} \right)^{\frac{1}{1-\alpha}} - \psi(\Pi_t - 1)\Pi_t Y_t + \psi\beta \mathbb{E}_t(\Pi_{t+1} - 1)\Pi_{t+1} Y_{t+1} \left(\frac{C_t}{C_{t+1}} \right)^\sigma = 0 \quad (46)$$

$$Y_t = A_t(\omega n_t^w + (1-\omega)n_t^p)^{1-\alpha} \quad (47)$$

$$1+i_t = \bar{i} + \left(\frac{\Pi_t}{\bar{\Pi}} \right)^{\delta_\pi} \left(\frac{Y_t}{\bar{Y}} \right)^{\delta_Y} \quad (48)$$

$$\frac{B_{t-1}}{\Pi_t} = \frac{B_t}{1+i_t} + (\omega T_t^w + (1-\omega)T_t^p) \quad (49)$$

$$T_t^w = T_t^p + z_t \quad (50)$$

$$B_t = \omega b_t^w + (1-\omega)b_t^p \quad (51)$$

$$Y_t = C_t + \frac{\psi}{2}(\Pi_t - 1)^2 Y_t \quad (52)$$

A.3 Calibration for Section 3

The parameters common to the attention trap and heterogeneous β models are set at standard values, shown in the first block of Table 2. Next is the novel parameter in the heterogeneous β model, the discount factor of poor households. The final portion of the table shows the parameter values for the information problem in the attention trap model. These are set such that in steady state $\beta^p(1+\bar{i}) = \beta(1+\bar{i}^p)$. The information cost has been set sufficiently low that q^w is close to 1 in all periods of the simulations (it remains within 10^{-40} of 1).

In all simulations, the model is solved nonlinearly using the perfect-foresight LMMCP solver for Dynare. A shock is realised in period 1, and from that point on all households and firms know the future transition paths of the economy.

It was initially estimated that the 2009 American Recovery and Reinvestment Act would add \$787 billion to US government debt (Congressional Budget Office, 2009), which amounts to 21.6% of the previous quarter's GDP. I set the shock in Sections 3.2.2 and 3.3.3 such that the ratio of government debt to previous quarter output rises by 7.2%

in the period that the shock hits in the baseline attention trap model.

Table 2: Calibration

Parameter	Description	Value
β	Discount factor	0.99
σ	Coefficient of relative risk aversion	1
φ	Disutility of labour	1
ω	Proportion of households which are wealthy	0.65
ϵ	Elasticity of substitution between intermediate good varieties	9
A	Steady state technology	1
α	Degree of diminishing returns in intermediate firms' production	0.25
ψ	Price adjustment costs	372.8
\mathcal{B}	Steady state government debt	1
$\bar{\Pi}$	Gross inflation target	1
d_{Π}	Elasticity of the interest rate to inflation	1.5
\bar{Y}	Output target	0.858954
\bar{i}	Interest rate target	$\beta^{-1} - 1 = 0.010101$
d_Y	Elasticity of the interest rate to output	0.125
β^p	Poor household discount factor	0.985
μ	Information cost	0.0001
g	Fraction of good financial institutions	0.15
τ	Inefficient bond dealer costs	$\frac{\beta - \beta^p}{\beta^p(1-g)(1+i)} = 0.005912$
τ^-	Lower inefficient bond dealer costs for Section 3.3.3	$\tau/2 = 0.002956$

B Proofs

B.1 Derivation of equation (3)

Suppose that there are N bond dealers, of which G are good. Each dealer is good with the same probability $g = \frac{G}{N}$. There are therefore S possible states of the world, where each state is a unique combination of matched dealers and interest rates, and each state is equally likely. S is given by:

$$S = \frac{N!}{G!(N-G)!} \quad (53)$$

Information processing is defined as the expected entropy reduction from prior to posterior beliefs, as is common in the rational inattention literature (Sims, 2003):

$$\mathcal{I} = H(h) - \mathbb{E}_s H(f|s) \quad (54)$$

Here H is the entropy function, a measure of the dispersion in a distribution:

$$H(\cdot) = - \sum_{k=1}^S \Pr(\text{state } k) \log \Pr(\text{state } k) \quad (55)$$

The distribution h is the prior belief, which is that each state is equally likely (i.e. no bond dealer is more likely to be good than any other). The distribution $f|s$ is the posterior, the distribution of beliefs implied by the signals households process, embodied in the decisions they took.

In the prior each state is equally likely, so:

$$H(h) = - \sum_{k=1}^S \frac{1}{S} \log \frac{1}{S} = \log S \quad (56)$$

Households choose signals that maximise their probability of successfully choosing a good bond dealer q for the minimum information cost. This means that if a signal instructs a household member to choose dealer d^* , their posterior belief will be that each state in which d^* is good is equally likely. There is no use in distinguishing between those states. The posterior therefore attaches probability $\frac{qN}{SG}$ to each of those states, where $\frac{SG}{N}$ is the number of states in which d^* is good. Similarly, they assign probability $\frac{(1-q)N}{S(N-G)}$ to each state of the world in which d^* is bad. The entropy of this posterior is therefore:

$$H(f|s \Rightarrow \text{choose } d^*) = -q \log \left(\frac{qN}{SG} \right) - (1-q) \log \left(\frac{(1-q)N}{S(N-G)} \right) \quad (57)$$

This is independent of the arbitrary d^* , so this expression also gives the expected entropy across all such signals. The total information processed to attain a success probability of q is therefore:

$$\mathcal{I} = \log S + q \log \left(\frac{qN}{SG} \right) + (1-q) \log \left(\frac{(1-q)N}{S(N-G)} \right) \quad (58)$$

Splitting apart the logs and cancelling terms we find:

$$\mathcal{I} = \log N - q \log G - (1-q) \log(N-G) + q \log q + (1-q) \log(1-q) \quad (59)$$

Now replace $G = gN$ and simplify to obtain equation (3):

$$\mathcal{I} = q \log \left(\frac{q}{g} \right) + (1-q) \log \left(\frac{1-q}{1-g} \right) \quad (60)$$

All terms involving the specific number of bond dealers N have cancelled out, leaving only terms in g , the proportion of bond dealers who are offering good interest rates. This

implies that as \mathcal{I} approaches $-\log(g)$, the success probability q approaches 1. While the information processing formula is not defined for $q = 1$, q gets arbitrarily close to 1 for this finite amount of information processing, which does not depend on the number of bond dealers in the market. $\mathcal{I}'(q) \rightarrow \infty$ as $q \rightarrow 1$.

B.2 $\mathcal{I}(\tilde{i}_t)$ is increasing and convex

Using the definition of the bad interest rate (1) we can write the effective interest rate in equation (2) as:

$$\frac{1}{1 + \tilde{i}_t} = \frac{1}{1 + i_t} + \tau(1 - q_t) \quad (61)$$

From this (dropping time subscripts and holding the policy rate constant) we have:

$$\frac{\partial q}{\partial \tilde{i}} = \frac{1}{\tau(1 + \tilde{i})^2} \quad (62)$$

From equation (3) we obtain:

$$\frac{\partial \mathcal{I}}{\partial q} = \log \left(\frac{q(1 - g)}{(1 - q)g} \right) \quad (63)$$

Combining these:

$$\frac{\partial \mathcal{I}}{\partial \tilde{i}} = \frac{\partial \mathcal{I}}{\partial q} \cdot \frac{\partial q}{\partial \tilde{i}} = \frac{1}{\tau(1 + \tilde{i})^2} \log \left(\frac{q(1 - g)}{(1 - q)g} \right) \geq 0 \quad (64)$$

We can guarantee that this is positive because $q \geq g$, so $\frac{q(1-g)}{(1-q)g} \geq 1$. Therefore $\mathcal{I}'(i_t^e)$ is weakly positive, and is strictly positive if $q > g$, so if $\mathcal{I} > 0$.

Differentiating this again with respect to \tilde{i} , we obtain:

$$\begin{aligned} \frac{\partial^2 \mathcal{I}}{\partial \tilde{i}^2} &= \frac{-2}{\tau(1 + \tilde{i})^3} \log \left(\frac{q(1 - g)}{(1 - q)g} \right) + \frac{1}{\tau(1 + \tilde{i})^2} \cdot \frac{\partial q}{\partial \tilde{i}} \cdot \frac{1}{q(1 - q)} \\ &= \frac{-2}{\tau(1 + \tilde{i})^3} \log \left(\frac{q(1 - g)}{(1 - q)g} \right) + \frac{1}{\tau^2(1 + \tilde{i})^4 q(1 - q)} \end{aligned} \quad (65)$$

This is > 0 iff:

$$\frac{1}{\tau(1 + \tilde{i})q(1 - q)} > 2 \log \left(\frac{q(1 - g)}{(1 - q)g} \right) \quad (66)$$

Substituting out for \tilde{i} using equation (2) this becomes:

$$\frac{1}{\tau(1 + i)q(1 - q)} + \frac{1}{q} > 2 \log \left(\frac{q(1 - g)}{(1 - q)g} \right) \quad (67)$$

$\mathcal{I}''(\tilde{i})$ is therefore positive if $\tau(1+i)$ is sufficiently low relative to g .

Taking exponentials of both sides and rearranging this becomes:

$$g > \frac{q}{(1-q)e^{Z(q)} + q}, \quad \text{where } Z(q) = \frac{1 + \tau(1+i)(1-q)}{2\tau(1+i)q(1-q)} \quad (68)$$

The right hand side of this condition is maximised (with respect to q) at:

$$q^* = \frac{-1 + \sqrt{1 + \tau(1+i) + \tau^2(1+i)^2}}{\tau(1+i)} \quad (69)$$

Substituting this in to equation (68) gives a lower bound on g in terms of τ and $(1+i)$ only:

$$g > \frac{\sqrt{K} - 1}{(1 - \sqrt{K} + \tau(1+i))e^{Z(q^*)} + \sqrt{K} - 1}, \quad \text{where } K = 1 + \tau(1+i) + \tau^2(1+i)^2 \quad (70)$$

The minimum g from condition 70 is positive, and increases with $\tau(1+i)$, but at any plausible level of $\tau(1+i)$ the restriction is extremely weak. Suppose, for example, that the interest rate on a good savings product is 20% a year, and the interest rate on a bad product is 0%. In a quarterly model, this implies an extremely high $\tau(1+i) = 0.047$, which implies that condition 70 becomes $g > 9 \times 10^{-20}$. For every one good bond dealer, there must be more than 1.1×10^{19} bad bond dealers before we approach the region where $\mathcal{I}''(\tilde{i}_t^e) > 0$ does not hold. Clearly, for plausible parameters, condition 70 will be satisfied.

B.3 The household first order conditions are sufficient for utility maximisation

Here I prove that for all plausible calibrations the household first order conditions (7 - 9) are sufficient for utility maximisation. I only consider a household away from the no-borrowing constraint, as at the constraint they fix attention at zero and their problem becomes identical to the poor household problem in existing TANK papers (e.g Bilbiie, 2008).

The utility maximisation problem of such a household can be written as an unconstrained maximisation by substituting out for consumption in the objective function (4) using the budget constraint (5):

$$\max_{n_t, b_t, \tilde{i}_t} U = \sum_t \beta^t \left(\frac{\left(\frac{b_{t-1}}{\Pi_t} + w_t n_t + \mathcal{D}_t - T_t - \frac{b_t}{1+i_t} \right)^{1-\sigma}}{1-\sigma} - \frac{n_t^{1+\varphi}}{1+\varphi} - \mu \mathcal{I}(\tilde{i}_t) \right) \quad (71)$$

The first order conditions are sufficient for utility maximisation if U is concave. This is

the case if the Hessian matrix is negative semi-definite, i.e. if for any real-valued vector $[x \ y \ z]$:

$$[x \ y \ z] \begin{bmatrix} \frac{\partial^2 U}{\partial \tilde{i}_t^2} & \frac{\partial^2 U}{\partial b_t \partial n_t} & \frac{\partial^2 U}{\partial b_t \partial \tilde{i}_t} \\ \frac{\partial^2 U}{\partial b_t \partial n_t} & \frac{\partial^2 U}{\partial n_t^2} & \frac{\partial^2 U}{\partial n_t \partial \tilde{i}_t} \\ \frac{\partial^2 U}{\partial b_t \partial \tilde{i}_t} & \frac{\partial^2 U}{\partial n_t \partial \tilde{i}_t} & \frac{\partial^2 U}{\partial \tilde{i}_t^2} \end{bmatrix} \begin{bmatrix} x \\ y \\ z \end{bmatrix} \leq 0 \quad (72)$$

Multiplying out the left hand side of this condition, substituting in the definitions of each second-order partial derivative (obtained using the unconstrained form of U in equation 71), and simplifying we obtain:

$$\begin{aligned} & -\sigma c_t^{-\sigma-1} \left(\frac{x}{1+\tilde{i}_t} - y w_t - \frac{z b_t}{(1+\tilde{i}_t)^2} \right)^2 - \varphi y^2 n_t^{\varphi-1} \\ & - \mathbb{E}_t \frac{\sigma \beta x^2 c_{t+1}^{-\sigma-1}}{\Pi_{t+1}^2} + \frac{2xz c_t^{-\sigma}}{(1+\tilde{i}_t)^2} - \frac{2z^2 b_t c_t^{-\sigma}}{(1+\tilde{i}_t)^3} - \mu z^2 \mathcal{I}''(\tilde{i}_t) \leq 0 \end{aligned} \quad (73)$$

This expression would clearly be true without the term in xz , which comes from $\frac{\partial^2 U}{\partial b_t \partial \tilde{i}_t}$. This is to be expected: it is the feedback between saving and attention (so interest rates) that is the cause for our concern that U may not be concave. I proceed by showing a condition under which this feedback is sufficiently weak that U remains concave, and so the first order conditions remain sufficient for utility maximisation. This condition is easily satisfied for plausible calibrations.

To derive the condition, I substitute out for the Euler equation (7) and first order condition on attention (9). When the resulting condition is satisfied, the first order conditions are therefore a local maximum of the utility function. However, since there is only one stationary point of the utility function and the function is continuous, this local maximum must be the global maximum.

We have from Appendix B.2 that:

$$\begin{aligned} \mathcal{I}''(\tilde{i}_t) &= \frac{-2}{\tau(1+\tilde{i}_t)^3} \log \left(\frac{q_t(1-g)}{(1-q_t)g} \right) + \frac{1}{\tau^2(1+\tilde{i}_t)^4 q_t(1-q_t)} \\ &= \frac{-2\mathcal{I}'(\tilde{i}_t)}{1+\tilde{i}_t} + \frac{1}{\tau^2(1+\tilde{i}_t)^4 q_t(1-q_t)} \end{aligned} \quad (74)$$

Substituting out for $\mathcal{I}'(\tilde{i}_t)$ using equation 9 gives:

$$\mathcal{I}''(\tilde{i}_t) = \frac{-2b_t c_t^{-\sigma}}{\mu(1+\tilde{i}_t)^3} + \frac{1}{\tau^2(1+\tilde{i}_t)^4 q_t(1-q_t)} \quad (75)$$

Substituting this in to condition 73, and also substituting out for $c_t^{-\sigma}$ using the Euler

equation (7), we have:

$$\begin{aligned}
& -\sigma c_t^{-\sigma-1} \left(\frac{x}{1+\tilde{i}_t} - yw_t - \frac{zb_t}{(1+\tilde{i}_t)^2} \right)^2 - \varphi y^2 n_t^{\varphi-1} \\
& - \mathbb{E}_t \frac{\sigma \beta x^2 c_{t+1}^{-\sigma-1}}{\Pi_{t+1}^2} + \mathbb{E}_t \frac{2\beta x z c_{t+1}^{-\sigma}}{\Pi_{t+1}(1+\tilde{i}_t)} - \frac{\mu z^2}{\tau^2(1+\tilde{i}_t)^4 q_t(1-q_t)} \leq 0 \quad (76)
\end{aligned}$$

A sufficient condition for this to be negative is if the sum of the final three terms is negative. After some algebraic manipulation this sufficient condition can be written as:

$$-\sigma \beta \mathbb{E}_t c_{t+1}^{-\sigma-1} \left(\frac{x}{\Pi_{t+1}} - \frac{z c_{t+1}}{\sigma(1+\tilde{i}_t)} \right)^2 + \mathbb{E}_t \frac{\beta z^2 c_{t+1}^{1-\sigma}}{\sigma(1+\tilde{i}_t)^2} - \frac{\mu z^2}{\tau^2(1+\tilde{i}_t)^4 q_t(1-q_t)} \leq 0 \quad (77)$$

Since the first term of this is always negative, the sufficient condition is satisfied if the sum of the final two terms is negative. That is, if:

$$\frac{\beta}{\sigma} \mathbb{E}_t c_{t+1}^{1-\sigma} \leq \frac{\mu}{\tau^2(1+\tilde{i}_t)^2 q_t(1-q_t)} \quad (78)$$

This is satisfied for all plausible calibrations. To see this, note that the right hand side of condition (78) is minimised with respect to q_t at:

$$q^* = \frac{1 + \tau(1 + i_t)}{2 + \tau(1 + i_t)} \quad (79)$$

Note that this is in terms of the policy rate i_t rather than the household effective rate \tilde{i}_t . At $q_t = q^*$, condition (78) becomes:

$$\frac{\beta}{\sigma} \mathbb{E}_t c_{t+1}^{1-\sigma} \leq \frac{4\mu(1 + \tau(1 + i_t))}{\tau^2(1 + i_t)^2} \quad (80)$$

With the calibration used in the simulations in Section 3, the left hand side of this always equals β . The quarterly policy rate i_t would have to exceed 240% for the condition to fail, and even this is a more restrictive condition than necessary, as condition 80 is merely a sufficient condition, and this calibration deliberately has a low value of μ to ensure that $q^w \approx 1$ in all periods. Alternatively, for the condition to fail at steady state (with $1 + i_t = \beta^{-1}$) we would require $\tau > 0.0201$, almost four times the value used in the simulations in this paper. Such high bad bond dealer costs would imply that while good products offer an annualised interest rate of 4% in steady state, bad products offer -4%. That is implausibly large for the interest rate dispersion on safe assets (see Macaulay (2020) for evidence for the UK).


C Intratemporal redistribution

Here I simulate a one-period redistribution shock z_t , which uses a rise in T_t^w to fund a cut in T_t^p . This is expansionary in both the attention trap and heterogeneous β models because poor households have a higher MPC than wealthy households.

As with debt-funded tax cuts, for small shocks poor households remain at their borrowing constraint and the two models are the same. If the shock is large enough that poor households temporarily save positive amounts, aggregate consumption responds less on impact, but is more persistent, in the attention trap model, for the same reasons discussed for debt-funded tax cuts in Section 3.

Using the same calibration as in Section 3, Figure 5a shows the response of aggregate consumption in the two models to such a redistribution. The shock is chosen such that poor households in the attention trap model accumulate the same stock of assets in period 1 as with the debt-funded tax cut considered in Section 3.2.2. Aggregate consumption rises by 48% less on impact in the attention trap model, but is more persistent, returning to steady state a quarter later than in the heterogeneous β model. Figure 5b shows why: effective interest rates facing poor households rise substantially relative to the policy rate in the attention trap model.

Figure 5: Effects of a redistributive tax shock



(a) Aggregate consumption in the attention trap model (solid) and the heterogeneous β model (dashed).

(b) Effective interest rates facing poor (solid) and wealthy (dashed) households in the attention trap model.

D Empirical exercises

D.1 Individual characteristics

Here I present some summary statistics for demographic variables across responsive and unresponsive states. There are no significant differences in race, gender, or education

composition before or after the tax cut, and the differences in age are quantitatively small, suggesting that changes in uncertainty between the groups before and after the tax cut are not caused by changes in the composition of the states.

Table 3: Demographic variables in the two groups of states. $k = 1$ indicates states whose GDP responds significantly to corporate tax cuts in the analysis of Liu and Williams (2019). The means post- tax cut are calculated using households entering the survey after the tax cut, so households in the survey both before and after the cut are counted in the pre- tax cut column.

Variable	Mean pre- tax cut		Difference	Mean post- tax cut		Difference
	$k = 0$	$k = 1$		$k = 0$	$k = 1$	
Age	49.70	50.46	-0.75** (-2.54)	49.63	50.89	-1.26** (-2.46)
Male	0.52	0.52	0.00 (0.02)	0.50	0.48	0.02 (0.99)
College degree	0.52	0.52	0.00 (0.20)	0.53	0.53	-0.00 (-0.21)
White	0.84	0.84	-0.00 (-0.58)	0.83	0.83	0.00 (0.17)

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note that all regressions in this paper except the final exercise in appendix D.2 contain individual fixed effects, which account for the differences in age and any other persistent individual characteristics. This means that the analysis concerns how uncertainty changes within the household as a result of the tax cut. This is the relevant question for assessing the empirical support for the model, because the key mechanism in the model is that tax cuts cause poor households to increase their savings, and therefore to increase their attention to saving choices. The mechanism therefore concerns changes within households.³⁷

D.2 Examining the mechanism

I have shown that uncertainty about savings interest rates falls after the 2018 tax cut in states likely to be responsive to that cut. To inspect the mechanism, I now examine how the tax cut affected two other variables. I first run regression 31 again, replacing

³⁷The evidence in Malmendier and Nagel (2011) (among others) suggests that households have persistently heterogeneous prior beliefs about assets. By controlling for individual fixed effects I do not allow such persistent belief heterogeneity to impact the results. In this way I focus on the mechanism in the model. Evidence of persistent unobserved heterogeneity can be seen in the fact that without individual fixed effects the correlation between log income and interest rate uncertainty is positive, while with fixed effects it is negative (though not significantly so), consistent with the model. Heterogeneous priors could explain this. If higher-income households have higher risk tolerances on average, they will save in riskier products, and so will have a higher prior uncertainty when asked about the return on savings products than a poorer household who will be thinking of less volatile products.

the dependent variable with the subjective probability of a rise in savings interest rates (from which the uncertainty in savings was computed).

Second, I examine the effects of the tax cut on the amount households intend to save, using the net contributions/withdrawals households expect to make to their savings accounts over the subsequent year. Since this comes from the annual household finance module of the survey, it is only available once per household, which rules out the use of individual fixed effects. I therefore run the alternative difference-in-difference model:

$$Y_{ist} = D_{st} \sum_{c=1}^4 \beta_c \mathbb{1}_{C_{ist}=c} + \gamma_c + \delta_s + \zeta_t + \theta X_{ist} + \varepsilon_{ist} \quad (81)$$

As with regression 31, the coefficients β_c give the estimated effect of the tax cut on the independent variable by income-fragility category. However, the individual fixed effects have been replaced by category fixed effects, and the vector of individual characteristics X_{ist} contains more variables, since it has to account for persistent individual heterogeneity due to the lack of individual fixed effects.³⁸ The results of the regressions on the probability of a rate rise and intended saving are in Table 4.³⁹

These checks are important. If the average reported probability of rates rising changed in responsive states relative to unresponsive states after the tax cut, the results in Table 1 could be picking up the effects of the tax cut on the expected level of the interest rate, rather than on uncertainty around the interest rate. The model predicts that the uncertainty fall for low income non-fragile households should have been driven by them starting to save more. The results of both regressions support the interpretation in section 4.

The first column of Table 4 shows that the tax cut did not have a significant effect on the subjective probability of a rise in interest rates. The fall in uncertainty associated with the tax cut found above is not driven by individuals getting substantially more (or less) optimistic on average about a rate rise. On average the views of individuals about whether rates will rise or fall remain unchanged, but in responsive states low income

³⁸I include age, age², gender, race, marital status, home ownership, employment status, education, tenure at current job, numerical ability score, and the number of times the respondent has been in the survey. I do not include income and financial fragility, since they are partially accounted for by the category dummies and there are already a large number of coefficients to estimate given the sample size. The treatment effect on the intended saving of low income non-fragile households is still significantly positive if these are included. I do not include differential time trends for treated and non-treated groups for the same reason. Treatment group-by-category time trends, if included, are all estimated to be insignificantly different from zero. Including them leads to a somewhat higher estimated treatment effect, suggesting that if anything pre-trends are biasing the treatment effect downwards.

³⁹For brevity I only report the results in the first column with robust standard errors. The treatment effects remain insignificant if movers are excluded and standard errors are clustered at the state level. In the regression on intended saving standard errors are clustered at the state level, as with only one observation per household there is no movement between states.

Table 4: Regression on interest rate forecasts and intended saving.

	(1)	(2)
	Probability of rate rise	Expected additions to savings
Low income non-fragile	-0.0126 (-1.20)	2369.7* (1.90)
High income non-fragile	0.00552 (0.49)	991.4 (0.46)
Low income fragile	0.00510 (0.49)	4067.9** (2.60)
High income fragile	-0.0128 (-1.02)	14679.3 (1.55)
Individual FE	Yes	No
State FE	Yes	Yes
Time FE	Yes	Yes
Clustered SE	No	Yes
Observations	102205	3368

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

non-fragile households move further away from a probability of 50%: they become more certain in their beliefs.

The second column shows that the tax cut is associated with a rise in intended saving for the group of interest. These two results are consistent with the model, in which tax cuts cause poorer households to pay more attention to their saving decisions, and so become less uncertain about savings markets, due to a rise in their savings.⁴⁰

⁴⁰An alternative test of the mechanism would consider the effects of the TCJA on incomes. However, what matters in the model is post-tax income, and the survey data reports pre-tax income. This is important since the TCJA had differential effects on income taxes across states, and indeed households in responsive states are also more likely to have experienced greater income tax cuts (using the state rankings calculated by Kumar (2020)). Pre-tax incomes in responsive and unresponsive states also have significantly different trends before the TCJA, which complicates the evaluation of treatment effects. I therefore leave the examination of how the TCJA affected pre-tax incomes to further research.