Overpersistence Bias in Individual Income Expectations and Its Aggregate Implications

Rozsypal, Filip; Schlafmann, Kathrin

Document Version
Accepted author manuscript

Published in:
American Economic Journal: Macroeconomics

DOI:
10.1257/mac.20190056

Publication date:
2023

Overpersistence Bias in Individual Income Expectations and its Aggregate Implications

By Filip Rozsypal and Kathrin Schlafmann*

Abstract

Using micro level data, we document systematic forecast errors in household income expectations that are related to the level of income. We show that these errors can be formalized by a modest deviation from rational expectations, where agents overestimate the persistence of their income process. We then investigate the implications of these distortions on consumption and savings behavior and find two effects. First, these distortions allow an otherwise fully optimization-based quantitative model to match the joint distribution of liquid assets and income. Second, the bias alters the distribution of marginal propensities to consume which makes government stimulus policies less effective.

JEL codes: D84, E21, D91
KEYWORDS: expectations, survey forecasts, savings, MPC

*Rozsypal: Danmarks Nationalbank and Centre for Macroeconomics (email: firo@nationalbanken.dk); Schlafmann: Copenhagen Business School, Institute for International Economic Studies, Danish Finance Institute (DFI) and Centre for Economic Policy Research (CEPR) (email: ksc.fi@cbs.dk). We would like to thank our discussant at NBER SI Chris Carroll for valuable comments that greatly improved the paper. We further would like to thank Vasco Carvalho, Wouter Den Haan, Chryssi Giannitsarou, Per Krusell, Hamish Low and Ricardo Reis as well as seminar participants at NBER SI Behavioral Macro, London School of Economics, Greater Stockholm Macro Group, University of Cambridge, University of Bonn, CEF Bordeaux, NorMac 2016, CEPR Household Finance 2016, Workshop on Household Surveys in Macroeconomics at University of Hamburg, FED Board, University of Essex, Danmarks Nationalbank, Deutsche Bundesbank, University of Illinois Urbana-Champaign, Brandeis University, University of Alberta, Norwegian Central Bank, University of Konstanz, Copenhagen Business School, Women in Macro and Finance Workshop Cologne, University of Copenhagen, E1Macro Workshop Queen Mary, Mannheim Quantitative Macro Workshop, Workshop on Expectation Formation Kiel, Uppsala University, European Midwest Micro and Macro Workshop 2018, NY Fed and UCL for many valuable comments and suggestions. Kathrin Schlafmann gratefully acknowledges financial support from Ragnar Söderbergs stiftelse and the Danish Finance Institute (DFI).

Fluctuations in income represent one of the most important sources of economic risk for households. Households who have different expectations about their future income realizations will hence make different decisions about consumption and saving today. Unfortunately, data on individual income expectations and corresponding realizations are not readily available. Despite the importance of household income expectations, testing their rationality or the identification of systematic biases has therefore been difficult.

In this paper we first use micro data on household income expectations, devise a new way to construct individual level forecasting errors and provide evidence of non-rationality in the form of a systematic bias related to the level of income. Second, we show that these empirical findings are consistent with a process of expectation formation where households are perfectly forward-looking but overestimate the persistence of their individual income process and are too pessimistic about the development of the aggregate economy. This formulation of expectations is in the spirit of Kahneman and Tversky (1973)'s finding of non-regression to the mean in people's probabilistic judgments. Third, we show how this bias affects consumption and savings behavior in an otherwise standard, fully optimization-based model of durable consumption. Including the bias allows the model to fit the joint distribution of liquid assets and income. In particular, this mechanism can explain why low income households do not borrow more to smooth consumption. We compare the model with biased income expectations to a rational model that requires a tighter borrowing constraint to fit the data. The two models generate different distributions of marginal propensities to consume, which results in government stimulus policies being less effective in the model with biased expectations: In a balanced budget experiment, the rational model predicts the increase in aggregate nondurable consumption to be 3.5 times larger than what the model with biased expectations predicts.

The first contribution of the paper is to empirically analyze forecast errors in individual household income expectations. Using data from the Michigan Surveys of Consumers, we show that current income is systematically correlated with the error people make when they forecast their individual future nominal income growth. Specifically, people in the upper part of the income distribution overestimate their future income growth while the opposite is true for lower income households: they are too pessimistic and underestimate their future income growth. In terms of magnitudes, on average people in the highest income quintile overestimate their income growth by 2 percentage points while people in the lowest income quintile underestimate it by 7 percentage points. Moreover, we show that people across the whole income distribution are too pessimistic about aggregate variables such as inflation and the unemployment rate. However, we show that these forecast errors about aggregate variables are not able to explain the magnitudes nor the differential signs of the forecast errors observed in individual nominal income growth expectations.

Analyzing errors in income expectations requires the knowledge of both a household's income expectation and the same household's ex post income realization over the corresponding time period. However, to the best of our knowledge none of the existing panel surveys satisfies these requirements. We exploit that the Michigan Surveys of Consumers reinterview a subset of households after 6 months. We obtain expectation errors for each household which allows us to document the systematic patterns in individual income forecast errors. We also ensure that our findings are not an artifact of the data construction procedure by conducting several robustness checks.

The second contribution of the paper is to present a rule for expectation formation that can explain the empirical findings and that is easy to implement in quantitative models. We show that the observed patterns in forecast errors are consistent with a form of expectation formation where people are fully forward looking but overestimate the persistence of their income process. We hence call this bias *overpersistence bias*. It implies that people overreact to shocks to their income and that this overreaction is persistent. The distorted expectations can be expressed as the sum of rational expectations and a function of current income, i.e. a function of all past shocks. Our formulation of expectations is therefore similar to "Diagostic Expectations" proposed by Gennaioli and Shleifer (2010) and Bordalo, Gennaioli and Shleifer (2018). The difference is that in their setup the bias term is a function of only the latest news, whereas in our setup it is a function of the full history of shocks.

The distorted expectations can be formulated parsimoniously in the context of a standard income process with persistent and transitory income shocks as in Storesletten, Telmer and Yaron (2004): We implement the overpersistence bias by allowing the agents' belief about the autocorrelation parameter to differ from the true underlying parameter. Moreover, we allow households to be too pessimistic about aggregate variables. This parsimonious representation of distorted expectations with only two free parameters is able to match the empirically observed forecast errors across the whole income distribution. The reason is that even though households share the same (distorted) beliefs about the data generating process of income, the overpersistence bias leads to heterogeneous expectation errors depending on the particular income realization of a given household. Households with currently high income expect their future income to remain higher than what their true income process would predict. Ex post they hence turn out to be too optimistic on average. The converse is true for households with currently low income: they underestimate their future income and turn out to be too pessimistic. While the overpersistence bias leads to heterogeneous effects, the aggregate pessimism affects households in the same way across the whole income distribution: People are too pessimistic about the future aggregate economy which biases downward their individual income expectations. We show that the combination of the two effects allows the expectations process to match the empirically observed magnitudes of forecast errors across the income distribution.

How does the overpersistence bias in income expectations affect the consumption-saving behavior of households? And what are the aggregate consequences of the bias? To answer these questions, we insert the fitted representation of expectation bias into an otherwise standard incomplete markets, heterogeneous agent model in the tradition of Bewley (1986) and Deaton (1991). Moreover, marginal propensities to consume (MPCs) have been the focus of much recent literature in the field of economics and household finance. Kaplan and Violante (2014) argue that it is crucial to include illiquid assets into the modeling framework to be able to capture the distribution of MPCs across the wealth distribution. In order to analyze how biased income expectations affect the distribution of MPCs we therefore include a durable consumption good in our analysis. We calibrate it to capture vehicle purchases as this allows the model to make simultaneous predictions about different categories of consumption which can be directly tested against empirical estimates.

Biased income expectations turn out to have differential effects on the behavior of households depending on their relative position in the income distribution. High income households hold similar portfolios under biased and under fully rational expectations. For them the overpersistence bias and aggregate pessimism have opposing effects and cancel each other out. In contrast, low income households choose different portfolios if they have biased income expectations. Low income households with biased expectations are too pessimistic about their future income and hence do not want to borrow to smooth consumption even though they would be able to borrow.

We show that this mechanism allows an otherwise standard, fully optimization-based model to fit the distribution of liquid assets as well as durable holdings across different income groups. In particular, including biased income expectations enables the model to match the distribution of liquid assets for low income households. The model with fully rational income expectations, on the other hand, would predict counterfactually large amounts of borrowing. Including the bias in income expectations as seen in the data allows the model to overcome this counterfactual behavior and to fit the distribution of borrowing.

Next, we study the effects of the bias on the level of MPCs for different consumption goods. Even though the MPCs were not targeted in the calibration, the model replicates the empirical estimates both in total consumption as well as the split between durable and nondurable goods: Households spend on average about 25% of the transfer on nondurable consumption in the first year. When taking into account durables expenditures this share increases to 75%. Moreover, the model also matches the heterogeneity in MPCs found by

Misra and Surico (2014): Only a small fraction of agents drives the large response in durable spending. We find that all else equal, biased agents have a smaller MPC in nondurable consumption than fully rational agents. However, they are more willing to shift their durable expenditures in time, resulting in higher overall marginal propensities to spend the transfer.

How much does the bias matter economically? We show that since the model with rational expectations generates larger amounts of borrowing than the model with biased expectation, it also requires a tighter borrowing constraint to fit the data. This tighter borrowing constraint amplifies the effects that the overpersistence bias has on MPCs. The rational model leads to a greater dispersion of MPCs across the income distribution: it generates a larger relative MPC of low income households compared to the MPC of high income households. The model with biased expectations, on the other hand, generates a relative MPC that is in line with empirical estimates (Johnson, Parker and Souleles, 2006; Parker et al., 2013). Relative MPCs are an important determinant of the government's ability to boost aggregate demand using fiscal transfers (Oh and Reis, 2012). Such policies have become popular during recessions. For example, the U.S. government handed out one-off cash transfers in both 2001 and 2008. To highlight the importance of the distribution of MPCs for these programs, we consider a balanced budget policy that levies a lump sum tax on high income households to pay for a lump sum transfer to low income households. In such an environment the aggregate consumption response is higher the larger is the difference between the MPC of households with high and low income. The differences in MPC distributions between the models with biased and rational expectations translate into a different assessment of the effectiveness of fiscal stimulus packages: The rational model predicts such policies to be much more effective than what the model with biased expectations predicts. In particular, the increase in aggregate nondurable consumption turns out to be 3.5 times larger than in the model with biased expectations.

The paper contributes to the literature in three fields. First, it contributes to the growing body of empirical studies analyzing expectations of households, firms and professional forecasters. To evaluate whether agents' expectations are rational one has to compare these expectations with the corresponding realizations. Most of this literature has therefore analyzed expectations about aggregate variables, where the realizations are readily available. Examples are Carroll (2003), Andolfatto, Hendry and Moran (2008), Malmendier and Nagel (2015), Coibion, Gorodnichenko and Kumar (2015), Cavallo, Cruces and Perez-Truglia (2017) and Vellekoop and Wiederholt (2017) for inflation expectations, Gerardi et al. (2008), Piazzesi and Schneider (2009), Case et al. (2012), and Kuchler and Zafar (2018) for house price expectations, Kuchler and Zafar (2018) for unemployment expectations, Piazzesi, Salomao and Schneider (2015) for expectations about excess bond returns and Bordalo, Gennaioli and

Shleifer (2018) for expectations about credit spreads. In contrast, we focus on *individual* level *income* expectations and realizations. Due to data availability, this area has received much less attention in the literature, Dominitz and Manski (1997), Dominitz (1998), Das and van Soest (1999) and recently D'Haultfoeuille, Gaillac and Maurel (2018) and Massenot and Pettinicchi (2018) being notable exceptions. Compared to the first two papers, the current paper has the advantage of analyzing a much larger sample of expectations and realizations, both in terms of the number of households and in terms of the time period covered. We are hence able to document systematic biases in household income expectations which are present throughout the past 25 years. Das and van Soest (1999) analyze household income expectations in a panel data set from the Netherlands, but in their data set households are only asked about the direction of expected income changes, not about the magnitude of these changes. While the authors also find that income expectations are too pessimistic in general they do not speak to the systematic bias we find with respect to the current level of income. Massenot and Pettinicchi (2018) also use data from the Netherlands and has similar data limitations. D'Haultfoeuille, Gaillac and Maurel (2018) propose an alternative test of rationality of expectations and use it to argue that income expectations are not rational. We build on Souleles (2004), who, using the same data set as the present paper, explored forecasting errors in a wide range of variables and noted the presence of systematic biases. We improve on his methodology of constructing the income forecast errors by explicitly taking the timing of survey questions into account. Studying the forecasting errors in a much more detailed way allows us to argue for overpersistence beliefs as the cause for the observed patterns in income expectation errors. The structural model further enables us to study the effects of this bias on savings and on the distribution of MPCs. Our paper is also related to a recent study by Das, Kuhnen and Nagel (2017). They document a relationship between socioeconomic status and expectations about a range of aggregate variables and interpret their results as low status agents being too pessimistic. In contrast, our findings suggest that all agents are too pessimistic towards aggregate outcomes. For individual income expectations, on the other hand, we find a differential effect: The overpersistence bias leads high income households to be too optimistic while low income households turn out to be too pessimistic.

One might worry about the degree to which agents act on the expectations they express in a survey. However, there is growing evidence that people indeed make decisions that are consistent with their stated beliefs. In a financially incentivized experiment in the context of inflation expectations, Armantier et al. (2015) show that agents' actions correlate with their beliefs. Moreover, there is growing evidence that agents' expectations are sensible in the sense that different ways of eliciting beliefs are consistent with each other (point forecast versus distributions) and actions – where available – are compatible with rational expectations

(De Bruin et al., 2011; Zafar, 2011). In an experiment in the housing context, Armona, Fuster and Zafar (2018) show that financial decisions of households are informed by their expectations. Wiswall and Zafar (2015) demonstrate the connection between expectations and actions in the context of education.

The second strand of literature this paper relates to is the formulation of expectation formation. Some of the recent research has focused on assessing whether predictable forecast errors – which are at odds with standard models of rational expectations – can be generated by rational models of information frictions such as sticky information (Mankiw and Reis, 2002) or noisy information (Woodford, 2003; Sims, 2003; Mackowiak and Wiederholt, 2009). Examples here include Coibion and Gorodnichenko (2012, 2015), Andrade and Le Bihan (2013) and Kohlhas and Walther (2018). On the other hand, an increasing number of studies suggest that decision makers do not form their expectations fully rationally (see, e.g., Cutler, Poterba and Summers (1990), DeLong et al. (1990), Greenwood and Shleifer (2014), Barberis et al. (2015), Gennaioli, Ma and Shleifer (2016), Fuhrer (2017), Barberis et al. (2018), Broer and Kohlhas (2018) and Carroll et al. (2018)). The paper that is the closest to the present study in this area is Bordalo, Gennaioli and Shleifer (2018). They propose that decision makers form their expectations under a representativeness bias, which effectively leads to overweighthing of the most recent innovation to income when forming expectations. In contrast, in the present setting where households overestimate the persistence of their income process, it is the current *level* of their income (rather than the last shock) that determines the forecasting error, which is supported by the predictive power of the level of income for the expectation errors that households make.

The third strand of literature that this paper directly contributes to is the literature on marginal propensities to consume. Empirically, examples for recent analyses include studies estimating the MPC out of government transfers (Johnson, Parker and Souleles, 2006; Parker et al., 2013; Misra and Surico, 2014; Parker, 2017), housing wealth (Mian, Rao and Sufi, 2013; Kaplan, Mitman and Violante, 2016), transitory income shocks (Jappelli and Pistaferri, 2014) and lottery winnings (Fagereng, Holm and Natvik, 2018). Recent structural models investigate the relationship between MPCs and wealth (see, e.g., Kaplan and Violante (2014) and Carroll et al. (2017)) and between MPCs and the business cycle (Berger and Vavra, 2015; Harmenberg and Öberg, 2017). The two most relevant studies for this paper in terms of modelling approach are Kaplan and Violante (2014) and Berger and Vavra (2015). Kaplan and Violante (2014) demonstrate that the presence of an asset with adjustment costs can generate realistic marginal propensities to consume out of transfer payments. Berger and Vavra (2015) show in a setting similar to ours that the phase of the business cycle further affects the MPC. We contribute to this literature by analyzing the effects of empirically relevant biases in income expectations on the behavior and MPC of households, in a model that matches both the empirical asset distributions as well as the MPCs in nondurable and durable consumption found in the literature. We show that biased and fully rational expectations have different implications for the joint distribution of liquid assets and income. Furthermore, the bias alters the distribution of MPCs across goods and across income, which affects the effectiveness of stimulus policies.

I Household Income Expectations in the Data

In this section, we analyze micro level data on household income expectations and show that low income households underestimate their income growth while high income households overestimate their income growth.

The data we analyze comes from the Michigan Surveys of Consumers. This survey interviews a representative cross-section of 500 households every month, with detailed expectation and income data available since July 1986. The households are asked about a wide range of topics, from expectations about the state of the aggregate economy, unemployment and inflation to purchasing conditions. Most importantly for the present analysis, people are also asked about their individual income expectations. Crucially, around one third of households are re-interviewed once after 6 months and they answer the same set of questions in both interviews. While we have income expectations for all households, for a subset of households we thus also have information about realized income growth.[1]

The survey asks households for their expected percentage growth in both income and prices. Specifically, the following questions are asked:

> *Q1a: During the next 12 months, do you expect your income to be higher or lower than during the past year?*
>
> *Q1b: By about what percent do you expect your income to (increase/decrease) during the next 12 months?*
>
> *Q2a: During the next 12 months, do you think that prices in general will go up, or go down, or stay where they are now?*
>
> *Q2b: By about what percent do you expect prices to go (up/down) on the average, during the next 12 months?*

[1] See the online appendix for a detailed description of the sample selection and a comparison of the income information with the Panel Study of Income Dynamics (PSID).

I.A Construction of Expectation Errors

The fact that a subsample of the surveyed households is re-interviewed after 6 months allows us to confront income growth expectations with realized income changes. The basic idea is to compare expected income growth with ex post realized income growth. The challenge is, however, that there is only imperfect overlap between the periods for which households give expectations and for which they report realizations. For our baseline analysis we therefore employ imputation methods to increase this overlap. To ensure that our results are neither driven by the imputation method nor by the imperfect overlap, we also conduct two robustness checks: First, we conduct the analysis on directly reported data for a subsample of households. This analysis is completely unaffected by imputation. Second, we analyze the subsample where after imputation the overlap is perfect. To the best of our knowledge, there is no other existing survey which would allow the direct comparison of income expectations and realizations.[2]

The exact data structure is as follows. When reporting their income, households are asked to state their total household income in the previous *calendar year*. Expectations, on the other hand, refer to *the following 12 months*. This has two implications. First, households who are interviewed for the first time in the first half of a year (January to June) report their income twice for the same time period since their re-interview falls into the same calendar year as the first interview. Households interviewed for the first time in the second half of a year (July to December), on the other hand, are re-interviewed in the next calendar year and hence report income for two consecutive years. Only for those households do we therefore have a reported income growth realization. Figure 1 illustrates the timing problem, showing as an example the data reported by households interviewed for the first time in January 2002 (panel (a)) and July 2002 (panel (b)), respectively. The second implication of the data structure, however, is that even for households interviewed in the second half of the year, the overlap between the reported income realizations and the time period that refers to the expectations is not perfect. Figure 1(b) shows that the overlap between expected and realized income is only 6 months for a household interviewed for the first time in July. This overlap is further decreasing for August to December households.

For our baseline analysis we exploit the fact that income growth reported by households interviewed in the second half of a year can be used to infer a relationship between this income growth in a particular year and the level of income as well as household characteristics in the year prior to that. We use this relationship to impute income growth realizations for the households interviewed in the first half of the year (see panel (c) of figure 1). Furthermore, to

[2]The online appendix contains a detailed discussion of income expectation and realization data in other surveys.

Figure 1: Timing of Income Realizations versus Expectations

(a) First interview in January 2002 - reported data:

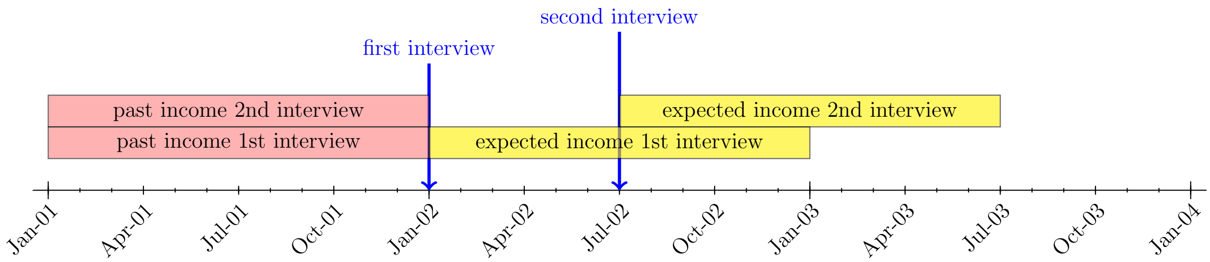

(b) First interview in July 2002 - reported data:

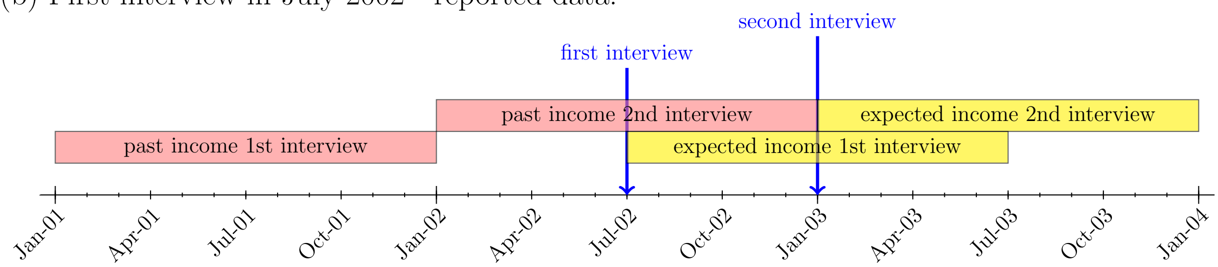

(c) First interview in January 2002 - imputed income:

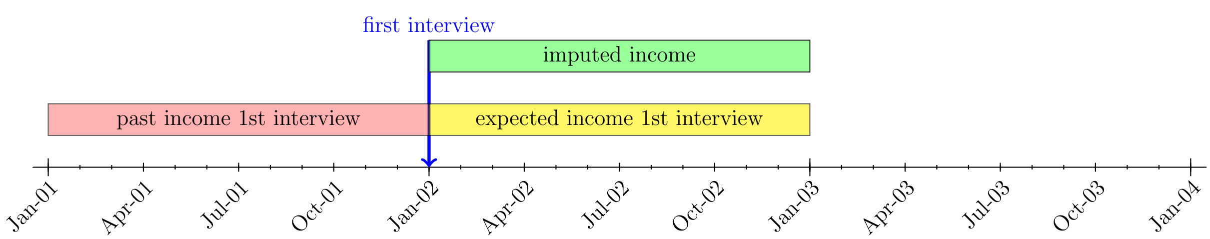

increase the overlap for households interviewed in the second half of the year, we impute their income growth using growth realizations of households interviewed in the following year. Imputation therefore both increases the number of observations and improves the timing overlap between expectations and realizations. We implement the imputation separately for each year. Our specification is therefore fully flexible regarding the effects of aggregate factors in the economy. A detailed description of the imputation procedure can be found in the online appendix.

To ensure that our findings are not an artifact of the imputation method, we conduct the analysis also on non-imputed data for July households as there is the largest overlap for

Figure 2: Expectation errors in real income growth

(a) mean (b) mean by income quintile

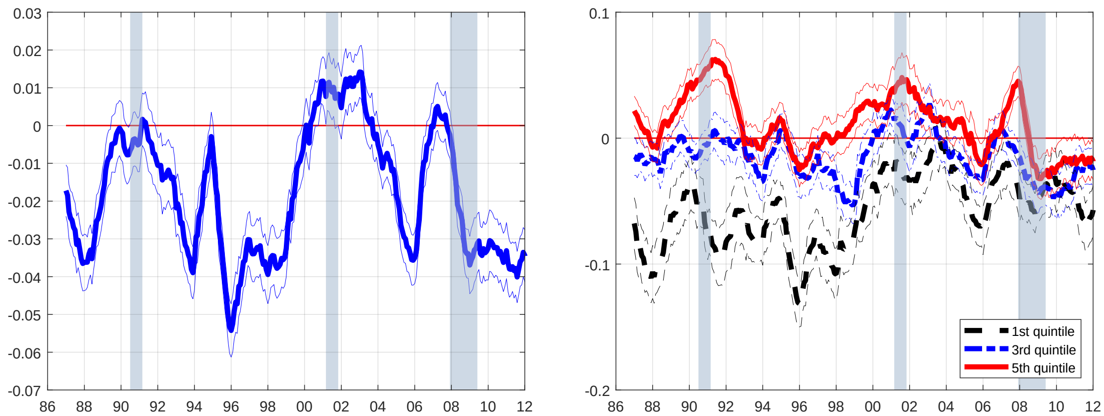

Note: The figure plots the mean expectation errors in individual real income growth smoothed with 12-month moving average filter. 95% confidence intervals (bootstrapped) are included. Expectation errors are winsorized at 5% and 95%. Households are allocated to income quintiles based on the cross-sectional distribution of per adult income in the year of the first interview. Data from the Michigan Surveys of Consumers and own calculations. Grey areas represent NBER recessions. On the y-axis, 0.01 corresponds to 1 percentage point.

directly reported data. Since we find similar results on this sample as we do on the full sample we can be assured that our results are not driven by the imputation procedure. Moreover, we conduct another robustness check to ensure that the imperfect timing of expectations and realizations does not affect our results. We re-run our analysis on the subsample of January, the month for which the timing overlap is perfect once we have imputed income growth realizations. Since our results also hold on this subsample we are confident the patterns we find are not driven by imperfect overlap of expectations and realizations either.

I.B Analysis of Expectation Errors

The expectation error of household i is constructed as

$$\psi_{i,t} = \hat{g}_{i,t+1|t} - \tilde{g}_{i,t+1}, \tag{1}$$

i.e. it is equal to the difference between the household's expected growth rate in income $\hat{g}_{i,t+1|t}$ and its realized growth rate $\tilde{g}_{i,t+1}$, where $\tilde{g}_i$ is either the imputed realized growth or the directly reported realized growth rate. Under this definition of the forecast error, a household who was too optimistic about its future income growth has a positive error.

Figure 2 shows the average expectation error in real income growth over the sample period.[3] For the population as a whole, people tend to be too pessimistic about their income growth (the average forecast error is mostly negative, see panel (a)). However, there is considerable heterogeneity in the forecast error by household income. While the low income group on average underestimates their income growth in all time periods, households in the high income group are in fact too optimistic for prolonged periods of time. Panel (b) shows the average expectation errors for three different income groups over time. Throughout the whole time period, the expectation errors are the lowest for the lowest income group (1st quintile) and highest for the highest income group (5th quintile).

Since households in different income quintiles are likely to also differ along other characteristics, we control for other observables using the following OLS regression:

$$Z_i = \alpha + \beta X_i + \sum_{k=1}^{K} \gamma_k D_{ik} + \varepsilon_i, \tag{2}$$

where Z_i is the outcome variable of interest of household i (in this case the expectation error ψ_i), X_i are household demographics as well as dummies for the month in which this household was interviewed, and D_{ik} are dummy variables which take the value 1 if household i belongs to income group k.[4] Table 1 shows the results of this regression. Even after controlling for other household characteristics, the effect of income in the first interview on expectation errors is highly significant and economically important. Looking at expectation errors in real income (column 1), households in the highest income quintile have on average an expectation error which is 3.5 percentage points more positive compared to households in the middle income group. At the same time, people in the lowest income group underestimate their income growth by 5.2 percentage points more than people in the middle income group.

Columns 2-4 repeat the analysis on different subsamples to ensure that the results are neither driven by imperfect overlap between the period of expectations and realizations nor by the imputation of realized changes. Columns 2 and 3 show the results when the sample is restricted to interviews in January or December only. For these months the overlap is perfect or almost perfect (11 out of 12 months), respectively. Since the results on these subsamples are very similar to the results on the full sample, we conclude that imperfect overlap does not generate our findings. Column 4 shows that the results also hold when the analysis is done on July interviews only using directly reported income changes instead

[3]In this section we focus our analysis on expectations about *real* income growth. However, the results we find are the same for *nominal* income expectations. Corresponding time series plots to figure 2 for nominal income expectations can be found in the online appendix. Moreover, when we control for household characteristics we will also show the regression results for errors in nominal income. These results will turn out to be very similar, both quantitatively and qualitatively, to the results for real income expectations.

[4]We have conducted many robustness checks to this specification, which can be found in the online appendix.

Table 1: OLS of expectation errors on household characteristics

	(1) real	(2) real	(3) real	(4) real	(5) nominal	(6) inflation
Income Quintile						
1 (low)	−0.052	−0.046	−0.049	−0.075	−0.049	0.004
	(0.006)	(0.018)	(0.027)	(0.021)	(0.007)	(0.000)
2	−0.018	−0.013	−0.025	−0.038	−0.016	0.002
	(0.006)	(0.017)	(0.024)	(0.020)	(0.006)	(0.000)
4	0.019	0.026	0.030	0.025	0.018	−0.002
	(0.005)	(0.013)	(0.024)	(0.016)	(0.005)	(0.000)
5 (high)	0.035	0.046	0.040	0.067	0.032	−0.004
	(0.006)	(0.015)	(0.022)	(0.017)	(0.006)	(0.000)
Education						
no high school	0.014	0.015	0.015	0.000	0.019	0.002
	(0.013)	(0.029)	(0.059)	(0.036)	(0.013)	(0.001)
college	−0.014	−0.024	−0.007	−0.032	−0.017	−0.003
	(0.004)	(0.012)	(0.016)	(0.013)	(0.004)	(0.000)
Age						
age	−0.004	−0.003	−0.007	−0.006	−0.004	0.000
	(0.001)	(0.003)	(0.006)	(0.004)	(0.002)	(0.000)
age × age	0.000	0.000	0.000	0.000	0.000	−0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Racial background						
black	0.019	0.025	0.009	0.021	0.024	0.002
	(0.008)	(0.018)	(0.032)	(0.022)	(0.008)	(0.000)
hispanic	0.013	0.005	0.018	0.018	0.018	0.003
	(0.009)	(0.027)	(0.046)	(0.033)	(0.009)	(0.001)
Number of adults						
1	−0.025	−0.004	−0.035	0.026	−0.025	0.001
	(0.009)	(0.026)	(0.039)	(0.042)	(0.010)	(0.001)
3 or more	0.020	0.014	0.021	0.021	0.018	−0.002
	(0.007)	(0.018)	(0.030)	(0.022)	(0.007)	(0.000)
Other family characteristics						
female	−0.008	−0.005	−0.007	−0.006	−0.002	0.005
	(0.004)	(0.010)	(0.016)	(0.012)	(0.004)	(0.000)
not married	0.023	0.004	0.030	−0.019	0.024	0.000
	(0.009)	(0.024)	(0.034)	(0.040)	(0.009)	(0.000)
Region						
North Central	−0.022	−0.023	−0.030	−0.020	−0.022	−0.000
	(0.006)	(0.015)	(0.024)	(0.017)	(0.006)	(0.000)
Northeast	−0.020	−0.021	−0.036	−0.005	−0.020	0.001
	(0.006)	(0.017)	(0.027)	(0.018)	(0.006)	(0.000)
South	−0.018	−0.014	−0.029	0.013	−0.017	0.001
	(0.006)	(0.016)	(0.024)	(0.016)	(0.006)	(0.000)
Constant	0.136	0.097	0.170	0.132	0.131	−0.016
	(0.052)	(0.078)	(0.148)	(0.094)	(0.054)	(0.002)
Sample	MAIN	JAN	DEC	JULY	MAIN	INF
Imputed Data?	yes	yes	yes	no	yes	no
Observations	58369	6973	2723	2805	58369	88017

Note: Regressions results from OLS of equation (2), where the dependent variable is the household expectation error in real income (columns 1-4), in nominal income (column 5) and in inflation (columns 6). The regressions include month dummies. Standard errors are reported in parentheses; they account for the uncertainty induced by imputation using multiple imputation procedures based on Rubin (1987), Barnard and Rubin (1999) and Reiter (2007).; without imputed data heteroskedasticity-robust standard errors are computed.

Figure 3: Expectation errors in real income by income group

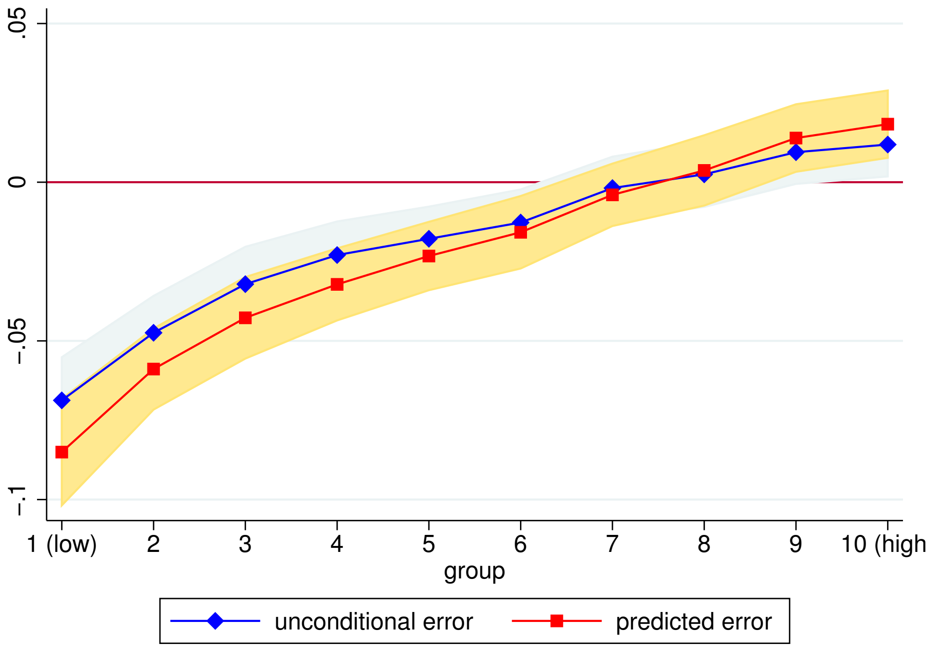

Note: The figure shows the unconditional mean expectation error (blue line, diamonds) and predicted expectation error (red line, squares) in real income growth by income decile. Predicted expectation errors are based on regression results from table 1 column 1, except that income is split in income deciles instead of quintiles. Predicted values are computed for all other explanatory variables at the weighted sample mean. Bands refer to 95% confidence intervals (standard errors account for the uncertainty that is induced by the imputation using multiple imputation procedures based on Rubin (1987), Barnard and Rubin (1999) and Reiter (2007).). On the y-axis, 0.05 corresponds to 5 percentage points.

of imputed ones. The sample in this specification is hence not affected by any imputation. The fact that the results hold confirms that the findings are not driven by the imputation procedure. Moreover, using a completely different method and data set, D'Haultfoeuille, Gaillac and Maurel (2018) also find low income households to be too pessimistic about their future income while high income households turn out to be too optimistic (see their section 6.2). Even though their short sample length does not allow them to control for aggregate effects, it is striking that they come to very similar conclusions to ours.

While the coefficients in table 1 are informative about the errors in the respective income group relative to the middle income group, they cannot directly tell us whether a particular income group is too optimistic or too pessimistic. Figure 3 thus plots both the unconditional mean expectation error by income decile and the expectation error predicted by the OLS regression when all other regressors are at their sample mean. The figure shows that while low income households underestimate their income growth, high income households are too optimistic and overestimate their income growth. In terms of magnitudes, on average people in the lowest income quintile underestimate their income growth by 7 percentage points and people in the highest income quintile overestimate it by 2 percentage points. The systematic relationship between forecast error and income group is thus robust to controlling for other household characteristics. In fact, as seen in figure 3, controlling for other demographics increases the effect of income on expectation bias.

Are households only systematically biased with respect to their individual income expectations? Or are they also biased in their expectations about aggregate conditions? In addition to the regression results for real income expectations, table 1 also splits the results in expectation errors in nominal income (column 5) and expectation errors in inflation (column 6). While income quintiles also have a significant effect on errors in inflation expectations, column 5 shows that most of the effects on expectation errors in real income are driven by the effects on expectation errors in nominal income. This is also confirmed in figure 4 where unconditional and predicted expectation errors are plotted for expectations in nominal income and inflation. The pattern for nominal income is very similar to that of real income. The reason for this small difference is that errors in inflation expectations are almost an order of magnitude smaller than errors in individual income expectations. Moreover, note that inflation expectations are too high across the whole income distribution. While there is an economically small variation in the size of errors in inflation expectations, this variation is not strong enough to change the sign of the bias as we move along the income distribution. The small impact of inflation expectations relative to income expectations is in line with Bachmann, Berg and Sims (2015) who find that consumers' spending attitudes are hardly affected by their inflation expectations.

Another aggregate variable that households in the Michigan Surveys of Consumers are asked about is unemployment. In particular, the question about unemployment expectations is the following:

> *How about people out of work during the coming 12 months – do you think that there will be more unemployment than now, about the same, or less?*

We code an expected increase in unemployment as -1, no change as 0 and and expected decrease as 1. This categorical expectation can be compared to the realized change in the U.S. unemployment rate in the 12 months following the interview. We code a realized change within +/- 0.1% as "0: no change", an increase in more than 0.1% as "-1: increase in unemployment" and a decrease of more than 0.1% as "+1: decrease in unemployment".[5] Categorical expectation errors are then defined as "categorical expectation" - "categorical

[5] We also computed all the analyses for alternative assumptions about the band for "the same" (+/- 0.05%, +/- 0.20% and +/- 0.25%) and the results were robust to these specifications.

Figure 4: Expectation errors by income group

(a) nominal income growth (b) inflation

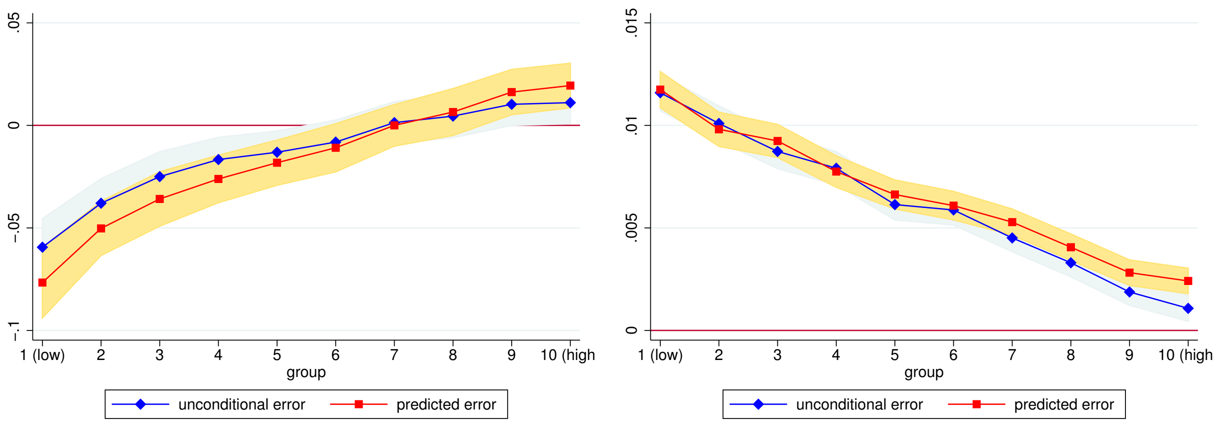

Note: The figure shows the unconditional mean expectation error (blue line, diamonds) and predicted expectation error (red line, squares) by income decile. Predicted expectation errors are based on regression results from table 1 column 5 and 6, except that income is split in income deciles instead of quintiles. Predicted values computed for all other explanatory variables at the weighted sample mean. Bands refer to 95% confidence intervals (for nominal income growth standard errors account for the uncertainty that is induced by the imputation using multiple imputation procedures based on Rubin (1987), Barnard and Rubin (1999) and Reiter (2007); for inflation heteroskedasticity-robust standard errors are computed). On the y-axis, 0.05 corresponds to 5 percentage points.

realization". The outcome categories for expectation errors range from "-2: far too pessimistic" to "+2: far too optimistic". We use an ordered logit regression to isolate the effect of individual income on errors in unemployment expectations (we keep the same control variables as in the analysis above).[6] Figure 5 shows the predicted likelihoods of each category for different income deciles, holding all other characteristics constant at their sample mean. The likelihood of a correct prediction is very stable around 55% to 58% for all income groups while the likelihood of being too pessimistic lies between 37% to 40%. At the same time, however, the likelihood of being too optimistic is very low for all income deciles. This indicates that - similarly to inflation expectations - people are too pessimistic across the whole income distribution. This finding of general pessimism in aggregate variables is in line with the results in Bhandari, Borovicka and Ho (2019) who show that unemployment and inflation expectations are on average too pessimistic across various population groups (including income groups) relative to the Survey of Professional Forecasters.

The analyses in this section thus reveal two forms of bias in household expectations. First, errors in individual income expectations vary systematically with income: Low income households underestimate their income growth while high income households overestimate their income growth. Second, households in all income groups are too pessimistic regarding

[6]Detailed regression results can be found in the online appendix.

Figure 5: Unemployment Expectations: predicted likelihood of each category by subgroups

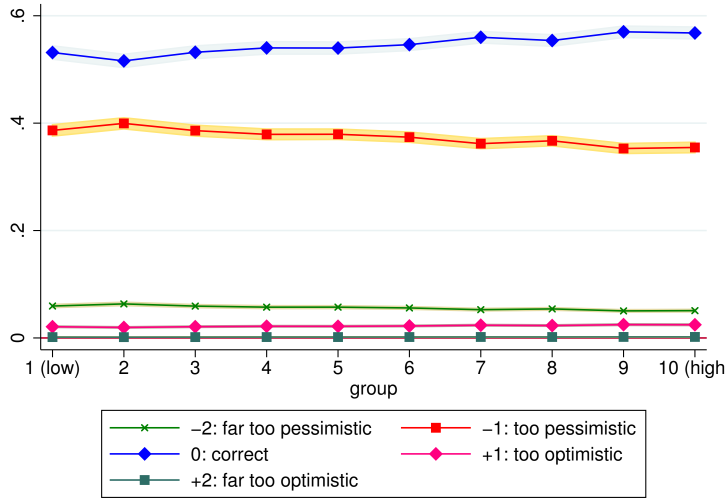

Note: The figure shows the predicted likelihoods of each outcome category of unemployment expectations (-2 (far too pessimistic) to +2 (far too optimistic)) by income decile. Predicted likelihoods are based on a ordered logit regression of categorical forecast errors on income deciles and other demographics as in previous regressions.

their forecasts of aggregate variables.

II Expectation Formation: Overestimation of Persistence in Income Process

In this section we present a formulation for expectation formation that can generate the observed pattern in expectation errors: We argue that people overestimate the persistence of their income process. This explanation can be seen as an expression of people's failure to properly account for regression to the mean in their probabilistic judgments (Kahneman and Tversky, 1973; Kahneman, 2012). While we cannot claim that this is the only mechanism that can generate the observed patterns, we did consider various alternative explanations and found that none of them was able to account for the observed joint distribution of income and expectation errors. A detailed description of the mechanisms considered and why they are not fully consistent with the observed data can be found in the online appendix.

II.A Mechanism: Overpersistence Bias

Formally, overestimating the persistence of income can be described as follows (for proofs of all results in this section see the online appendix). Assume that income (net of age effects and the effects of other demographics) is generated by the process

$$\ln Y_{i,t} = \ln P_{i,t} + \ln T_{i,t}, \tag{3}$$

$$\ln P_{i,t} = \rho \ln P_{i,t-1} + \ln N_{i,t}, \tag{4}$$

where P_{it} is a persistent component and T_{it} is a transitory shock. Persistent income depends on past persistent income and on a shock N_{it}. Both shocks are independently and log-normally distributed with mean 1. Overestimating the persistence implies that the households believe their persistence parameter to be larger than it actually is:

$$1 > \hat{\rho} > \rho \tag{5}$$

Theorem *If the true income process is governed by equations* (3) *and* (4) *and the household overestimates the persistence of the process according to equation* (5),

(a) $\exists! \bar{P}$:

$$\mathsf{E}\left[\hat{\mathsf{E}}_t[\ln(Y_{i,t+1})] - \ln(Y_{i,t+1}) | P_{i,t} > \bar{P}\right] > 0$$

and vice versa for $P_{it} < \bar{P}$, *where* $\hat{\mathsf{E}}_t[\ln(Y_{i,t+1})]$ *is the distorted expectation of* $Y_{i,t+1}$ *given information at time* t.

(b) *Let* $\Delta_{i,t} \equiv P_{i,t} - \bar{P}$, *then*

$$\frac{\partial \mathsf{E}\left[\hat{\mathsf{E}}_t[\ln(Y_{i,t+1})] - \ln(Y_{i,t+1}) | \Delta_{i,t}\right]}{\partial \Delta_{i,t}} > 0.$$

The proposition thus states that overestimating the persistence of the income process generates expectation errors in income growth that are (a) positive if the persistent income component is above a certain threshold (and negative if it is below this threshold) and (b) increasing in the distance from this threshold. Overpersistence can hence generate the pattern of systematic expectation errors observed in figure 3.

Intuitively, overestimating the persistence of the income process has the effect that people do not sufficiently account for mean-reversion of income in the cross-section. This interpretation is supported by figure 6. Panel (a) shows that income is indeed mean-reverting by plotting the realized real income growth rates that are predicted for each income decile if all other household characteristics are at their sample mean. Low income households are predicted to experience a large income growth and the predicted growth is decreasing in income. High income households, in fact, are predicted to have a negative income growth.

Figure 6: Realized growth and growth expectations in real income by income group

(a) realized growth (b) growth expectation

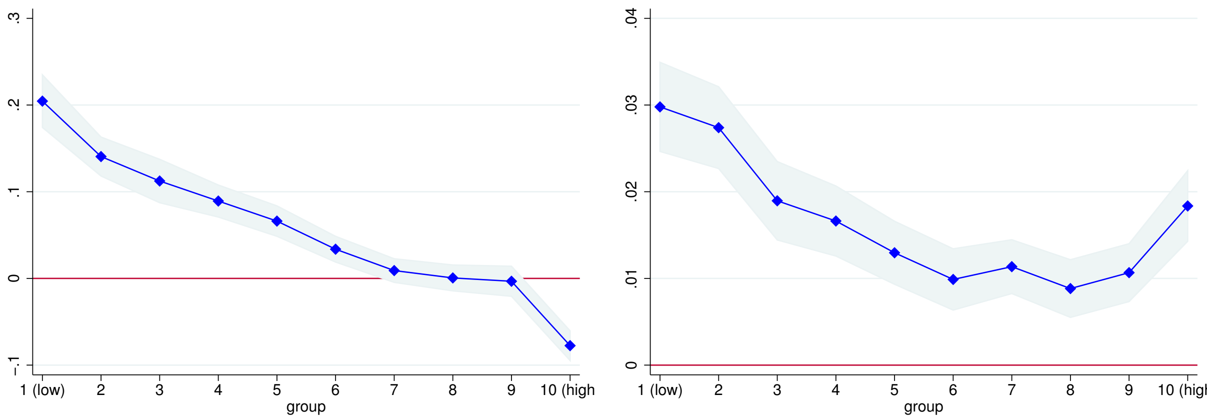

Note: The figure shows the predicted realized growth (panel (a)) and growth expectations (panel (b)) in real income by income decile. Predicted values are based on OLS regression results from regressing individual realized growth rates or expectations on all regressors as in table 1. Detailed estimation results can be found in the online appendix. Sample: for realized growth only directly reported income growth rates are used (first interviews in second half of the year); for growth expectations all observations are used (with or without reinterview and all months). Predicted values computed for all other explanatory variables at the weighted sample mean. Bands refer to 95% confidence intervals (based on heteroskedasticity-robust standard errors). On the y-axis, 0.01 corresponds to 1 percentage point.

Panel (b) further plots the growth expectations that are predicted for each income decile, again holding all other characteristics constant at their sample mean. Growth expectations, like realized income growth, decrease with income. However, comparing the magnitudes we see that households fail to anticipate the magnitude of the mean reversion. We interpret this finding as evidence in favor of households overestimating the persistence of their income process.

The expectations under the overpersistence bias can also be expressed as a function of rational expectations and the history of past innovations:

Corollary *If the true income process is governed by equations* (3) *and* (4) *and the household overestimates the persistence of the process according to equation* (5), *the distorted expectation at time t of income in period $t = t+1$, $\hat{\mathsf{E}}_t[\ln Y_{i,t+1}] = \hat{\rho} \ln P_t$, can be expressed as*

$$(6) \qquad \hat{\mathsf{E}}_t[\ln Y_{i,t+1}] = \mathsf{E}_t[\ln Y_{i,t+1}] + (\hat{\rho} - \rho) \cdot \sum_{s=0}^{\infty} \rho^{s-1} \big(\mathsf{E}_{t-s}[\ln Y_{i,t-s+1}] - \mathsf{E}_{t-s-1}[\ln Y_{i,t-s+1}] \big)$$

where $\mathsf{E}_t[\ln Y_{i,t+1}] = \rho \ln P_t$ is the rational expectation of income in period $t+1$ based on information available at time t.

This implies that due to the overpersistence bias the distorted beliefs are equal to the sum of the rational expectation and a weighted sum of all innovations to past rational expectations. People under the overpersistence bias hence overreact to income shocks and the overreaction to a specific shock is persistent but decaying over time. This formulation of expectation formation is related to expectations formed by "Diagnostic Expectations" proposed in Gennaioli and Shleifer (2010) and Bordalo, Gennaioli and Shleifer (2018).[7] The difference is that in their setup the distortion would only be a function of the latest shock, $\hat{\mathsf{E}}_t[\ln Y_{i,t+1}] = \mathsf{E}_t[\ln Y_{i,t+1}] + \theta \cdot \left(\mathsf{E}_t[\ln Y_{i,t+1}] - \mathsf{E}_{t-1}[\ln Y_{i,t+1}]\right)$, where the parameter θ governs the magnitude of the bias due to diagnostic expectations. In contrast, with the overpersistence bias distortions accumulate over time. This persistence in distortions explains why empirically the level of income is systematically related to the forecast error households make.

II.B Modeling and Quantifying Biased Beliefs

From the analyses in the previous sections we conclude that there are two forms of systematic bias in household income expectations: First, low income households are too pessimistic about their income growth while high income households are too optimistic. This pattern is consistent with people overestimating the persistence of their income process. Second, households across the whole income distribution are too pessimistic about aggregate conditions. We will now formulate how to parsimoniously incorporate these distortions in a model framework and quantify their magnitudes by matching the expectation errors in the model with those documented in the data.

We proceed in three steps. First, we assume a particular type of income process that is typically used in the quantitative literature (see, e.g., Berger and Vavra (2015) and Storesletten, Telmer and Yaron (2004)) and parametrize this process using standard estimates from the literature. Second, we allow households to have wrong beliefs about the persistence of the process as well as to be too pessimistic about aggregate developments. Third, we calibrate these two belief parameters and show that this parsimonious representation is able to

[7] Note that mathematically, the overpersistence bias can be expressed in the general framework of Bordalo, Gennaioli and Shleifer (2018):

$$h^{\theta}(\ln \hat{P}_{i,t+1}) = h(\ln \hat{P}_{i,t+1} | \ln P_{i,t} = \ln \hat{P}_{i,t}) \cdot \left(\frac{h(\ln \hat{P}_{i,t+1} | \ln P_{i,t} = \ln \hat{P}_{i,t})}{h(\ln \hat{P}_{i,t+1} | \ln P_{i,t} = (\rho - 1) \ln \hat{P}_{i,t})} \right)^{\theta} \frac{1}{Z}$$

where $h^{\theta}(\ln \hat{P}_{i,t+1})$ is the distorted probability distribution, $\theta = \hat{\rho} - \rho$, $h(\ln \hat{P}_{i,t+1} | \ln P_{i,t} = \ln \hat{P}_{i,t})$ is the true probability distribution based on current information and $h(\ln \hat{P}_{i,t+1} | \ln P_{i,t} = (\rho - 1) \ln \hat{P}_{i,t})$ is a specific reference distribution, which in this case is a normal distribution with mean $(\rho - 1) \ln \hat{P}_{i,t}$ and variance $var(\ln N_{i,t})$. This is a different reference distribution compared to the one Bordalo, Gennaioli and Shleifer (2018) employ in their paper.

replicate the observed expectations errors across the income distribution.

Underlying Income Process The exogenous income of a household is a combination of three mutually independent exogenous components: a persistent aggregate component Z_t, a persistent idiosyncratic component $P_{i,t}$ and a idiosyncratic transitory component $T_{i,t}$:

$$Y_{i,t} = Z_t \cdot P_{i,t} \cdot T_{i,t}. \tag{7}$$

Transitory shocks $T_{i,t}$ are iid lognormally distributed with

$$T_{i,t} \sim \log N\left(-\sigma_T^2/2, \sigma_T^2\right). \tag{8}$$

The idiosyncratic persistent component $P_{i,t}$ follows an AR(1) process in logs such that

$$\ln P_{i,t} = \rho \ln P_{i,t-1} + \epsilon_{i,t}^P, \qquad \epsilon_{i,t}^P \sim N(0, \sigma_P^2) \tag{9}$$

and the aggregate persistent component is a two state Markov process

$$\mathbb{Z} = \begin{bmatrix} Z^h \\ Z^l \end{bmatrix}, \qquad \Pi_Z = \begin{bmatrix} \pi_{11} & 1-\pi_{11} \\ 1-\pi_{22} & \pi_{22} \end{bmatrix}, \tag{10}$$

where the high state refers to boom periods and the low state to recessions.

Incorporating Beliefs Motivated by our findings discussed above, we allow households to have biased beliefs about their income process. The overpersistence bias in expectations is implemented by allowing agents to believe that the persistence of the idiosyncratic component P is different than its true value. Formally, agents believe that their persistent income component evolves according to the following process:

$$\ln P_{i,t} = \hat{\rho} \ln P_{i,t-1} + \epsilon_{i,t}^P, \qquad \epsilon_{i,t}^P \sim N(0, \sigma_P^2), \tag{11}$$

where the persistence belief $\hat{\rho}$ is allowed to differ from the true persistence of the process ρ.

The pessimism in aggregate developments is implemented by allowing agents to believe that the level of the aggregate states will differ from the true levels by a factor μ:

$$\hat{Z}_{t+1} = \mu \mathsf{E} Z_{t+1} = \mu \Pi_Z(Z_t)\mathbb{Z}, \tag{12}$$

where $\Pi_Z(Z_t)$ is the row of Π_Z that corresponds to Z_t. To quantify the biases, we find both bias parameters - the overpersistence belief $\hat{\rho}$ and the pessimism parameter μ - by matching the empirically observed forecasting errors by income quintile with the ones generated in this model.

Matching Expectation Errors Before fitting the bias parameters we need to parametrize the true income process. In the literature there is a debate about the true persistence of household income. Here in the main text we follow Storesletten, Telmer and Yaron (2004) who estimate an income process with persistent and idiosyncratic shocks. In the online appendix we show that this choice is not crucial: the overpersistence bias is able to match the observed forecast errors also for higher values of the persistence parameter, including the limit of a random walk. We transform Storesletten, Telmer and Yaron (2004)'s income process to quarterly frequency and obtain the following parameters: The persistent income component has an autocorrelation parameter of $\rho = 0.9774$ with standard deviation $\sigma_P = 0.0424$. The transitory income shocks have a standard deviation of $\sigma_T = 0.1$. To determine the transition matrix for the aggregate component of income we target the average duration of NBER recessions and booms in the post-war period (1945-2009).[8] On average in this period, booms lasted 58.4 months while recessions lasted 11.1 months. This leads to the probability of entering a recession of 6.85% and of leaving a recession of 36.04%. The levels of the boom and recession states have been chosen to reflect the average positive and the average negative deviation from trend in HP-filtered GDP. The resulting levels of booms and recessions are 1.0040 and 0.9790, respectively.[9]

We choose the overpersistence parameter $\hat{\rho}$ and the aggregate pessimism parameter μ to match the empirically observed expectation errors by income group. The parameters that match the errors are $\hat{\rho} = 0.9831$ (compared to the true persistence of $\rho = 0.9774$) and $\mu = 0.9778$. Table 2 shows that with these two parameters the model is able to match the expectation errors for all five income quintiles perfectly up to the second digit: The overpersistence belief generates the spread across the income distribution while the aggregate pessimism shifts down the expectations errors for all income groups.

Another benefit of the parsimony of this specification is that it makes the bias simple to implement in various settings. In the remainder of this paper, we focus on consumption-

[8]This specification leads to an asymmetric transition matrix. As a robustness check we have run all analyses (both the quantification of the biases as well as the solution of the complete model of consumption in the next section) also with a symmetric specification where we let the aggregate component Z_t follow an AR(1) process, parametrized as in Berger and Vavra (2015). Under this specification, all the results remain qualitatively identical and quantitatively very similar.

[9]The exact formula is

$$\text{avg\_dev} = \frac{1}{T_{pos}} \sum_{t=1}^{T} \hat{y}_t \cdot I(\hat{y}_t > 0) - \frac{1}{T_{neg}} \sum_{t=1}^{T} \hat{y}_t \cdot I(\hat{y}_t < 0)$$

where T_{pos} (T_{neg}) is the number of periods where $\hat{y}$ is *positive* (*negative*) in the sample and $\hat{y}_t$ is HP-filtered log(GDP). This difference between the good and the bad state combined with the fraction of time spent in booms and recessions (which results from the transition matrix) as well as the constraint that the mean of the overall process is 1 gives the levels of the two states.

Table 2: Mean expectation errors

	data	model
income quintile 1	-0.072	-0.069
income quintile 2	-0.037	-0.037
income quintile 3	-0.019	-0.021
income quintile 4	-0.000	-0.007
income quintile 5	0.016	0.021

Note: Data moments are the expectation errors predicted by equation (2) when all control variables apart from income are held constant at their sample mean.

saving implications. However, using this specification it would be straightforward to implement and study the overpersistence bias in other settings, for example in a model of asset pricing.

III Implications of Biased Income Expectations

In this section we analyze how the distortions that we documented in income expectations affect consumption and saving decisions and investigate their aggregate implications. To do so we insert the representation of beliefs that we fitted in the previous section into a standard incomplete markets, heterogeneous agent model in the tradition of Bewley (1986) and Deaton (1991). Kaplan and Violante (2014) argue that it is crucial to include an illiquid asset into structural models to be able to match MPCs across the wealth distribution. To be able to meaningfully analyze the distribution of MPCs we therefore include a durable good into our quantitative analysis. Our model setting is close to the one used by Berger and Vavra (2015). Apart from allowing for biased income expectations the most important difference is in the treatment of the borrowing constraint. Whereas Berger and Vavra (2015) assume that agents can only save (no borrowing), we allow households to borrow up to a limit determined by their income state and durable holdings. This assumption is not only more realistic, but it also has important consequences. A significant fraction of US households holds negative liquid assets. In order for the model to fit the data borrowing is hence essential. However, fitting the distribution of how much people borrow, as opposed to only the fraction of households that borrow, is challenging for the class of models that we study. In section 4.2, we show that including the bias in income expectations as seen in the data allows the model to replicate the empirical distribution of borrowing. In section 4.3, we use the calibrated model to analyse how the overpersistence bias alters the behaviour compared to an identically parameterized model under rational expectations. Finally, in section 4.4, we demonstrate the economic importance of the bias. We turn to a simplified calibration where we calibrate an exogenous borrowing constraint to fit the share of households with positive liquid assets. Since the rational model generates more borrowing, it requires a tighter borrowing constraint to fit this data moment. We show that this amplifies the effects of biased expectations, which leads to economically large differences in the distribution of MPCs and hence to large differences in the assessment of government stimulus policies.

III.A Model Setup

We consider the following partial equilibrium framework. Households are infinitely lived and derive utility from two sources: a non-durable consumption good and a flow of services from a durable good. The stock of durable goods depreciates and is subject to non-convex adjustment costs. Households hence optimally adjust their durable holdings only infrequently. In addition to durable goods, households can also invest in a riskless liquid asset which they can also use to borrow. The only source of risk the households face are fluctuations in their exogenous income.

Households maximize their discounted life time utility (to simplify notation we have dropped the subscript i which indicates the individual household)

$$\max_{\{c_t\}_{t=0}^{\infty},\{d_t\}_{t=0}^{\infty},\{s_t\}_{t=0}^{\infty}} \quad \sum_{t=0}^{\infty} \beta^t \ \mathsf{E}\big[U(c_t, d_t)\big], \tag{13}$$

subject to the following budget constraint

$$c_t + d_t + s_t + A(d_t, d_{t-1}) \leq Y_t + (1-\delta)d_{t-1} + R(s_{t-1}). \tag{14}$$

Households have available resources based on their income Y_t, the value of their depreciated durable stock $(1-\delta)d_{t-1}$, and the current value of the liquid asset holdings they chose in the previous period $R(s_{t-1})$. The current value of their liquid assets is determined as follows:

$$R(s_t) = [1 + r\,(s_t)]s_t \text{ where } r(s_t) = \begin{cases} r^l & \text{if } \ s_t > 0 \\ r^b & \text{if } \ -(\kappa_y P_t + \kappa_v d_t) \leq s_t \leq 0 \end{cases} \tag{15}$$

where $r^b > r^l$. Households can either save or borrow in liquid assets but have to pay a higher rate of interest for borrowing than they obtain when they are saving. The borrowing limit $(\kappa_y P_t + \kappa_v d_t)$ depends on their current persistent income (a loan-to-income constraint $\kappa_y P_t$) and the value of their durable stock (a loan-to-value constraint $\kappa_v d_t$). Our endogenous specification of the borrowing constraint departs from the practice of a fixed borrowing limit that is prevalent in the literature. We will show in section III.D that under the assumption of a fixed borrowing limit, the model with biased income expectations requires a less restrictive borrowing limit to fit the data than a rational model. This turns out to have significant consequences for the distribution of marginal propensities to consume and for the effectiveness of stimulus policies.

Households spend their available resources on non-durable consumption c_t, liquid assets s_t and the new durable stock d_t subject to adjustment costs $A(d_t, d_{t-1})$:

$$A(d_t, d_{t-1}) = \begin{cases} 0 & \text{if } d_t = (1-\delta)d_{t-1} \\ F^d(1-\delta)d_{t-1} & \text{otherwise.} \end{cases} \tag{16}$$

Equation (16) states that there are no adjustment costs if the household chooses to keep its depreciated durable stock, i.e. $d_t = (1-\delta)d_{t-1}$. On the other hand, if the household adjusts its durable stock, it has to pay adjustment costs equal to fraction F^d of the depreciated stock before the it is free to choose any new level of durable stock d_t.

Finally, the period utility function is

$$U(c,d) = \frac{\left[\left((1-\theta)c^{\frac{\xi-1}{\xi}} + \theta(\bar{d}+d)^{\frac{\xi-1}{\xi}}\right)^{\frac{\xi}{\xi-1}}\right]^{1-\gamma}}{1-\gamma}. \tag{17}$$

Note that every household obtains utility from a small free stock of durable $\bar{d}$. This captures the fact that even a very old car with almost zero resale value can be used as means of transport. This specification of the utility function hence enables the model to match the empirical distribution of durable stocks with its substantial share of low values.

The only source of risk in the model is income risk. We assume that income follows the process as described in the previous section (equations (7)-(10)) and that households have biased beliefs according to equations (11) and (12).

III.B Matching the Model to the Data

The model is calibrated at quarterly frequency. We proceed in two steps. First, we set the parameters of the environment (interest rates, borrowing constraints, depreciation rate and adjustment costs) exogenously according to either our empirical estimates or results from the literature. Second, we calibrate the remaining preference parameters to match the empirical distributions of liquid assets and durable holdings. Note that the belief parameters are independent of the specification of the consumption model so that we can use the parameters obtained in the previous section. Table 3 reports the complete parametrization.

Exogenous Parameters of the Environment Households can both save and borrow in the liquid asset but earn a rate of return that depends on their balance. The interest rate for saving is set to the mean real interest rate on 3 month treasury bills in the post-war

Table 3: Parameter Values

Parameter		Value
technology:		
interest rate (lending)	r^l	0.0016
interest rate (borrowing)	r^b	0.02
loan-to-income constraint	κ_y	0.56
loan-to-value constraint	κ_v	0.8
depreciation rate	δ	0.05
adjustment costs	F^d	0.3
income:		
persistence of idiosyncratic income process	ρ	0.9774
std dev of idiosyncratic persistent shocks	σ_P	0.0424
std dev of idiosyncratic transitory shocks	σ_T	0.1
high aggregate income state	Z^h	1.0040
low aggregate income state	Z^l	0.9790
prob. of entering recession	$1-\pi_{11}$	6.85%
prob. of leaving recession	$1-\pi_{22}$	36.04%
beliefs:		
persistence of income	$\hat{\rho}$	0.9831
aggregate pessimism	μ	0.9778
preferences:		
discount factor	β	0.9825
risk aversion	γ	1.5
weight of durable goods in utility	θ	0.075
elasticity of substitution in utility	ξ	3
free durable services	$\bar{d}$	0.5

period (1948-2015). On quarterly frequency this value is equal to $r^l = 0.0016$. The interest rate for borrowing is set equal to $r^b = 0.02$ which reflects interest rates on credit cards and on auto loans. Data on credit card rates is available since 1994 ("Commercial Bank Interest Rate on Credit Card Plans, All Accounts") and interest rates on auto loans since 1972 ("Finance Rate on Consumer Installment Loans at Commercial Banks, New Autos 48 Month Loan"). The mean real interest rates on quarterly frequency for these two series are 0.0268 and 0.0127, respectively. Since households in the model borrow at the same rate against their income (which reflects credit card debt) and against durables (which resembles auto loans), we set the borrowing rate to 0.02, a value that is roughly in the middle of the two interest rates. Moreover, this value is well within the range of interest rates on car loans for new and used cars documented by Attanasio, Goldberg and Kyriazidou (2008) for the Consumer Expenditure Survey.

To set the loan-to-income constraint we turn to data from the Survey of Consumer

Figure 7: Model fit

(a) distribution of durable goods (b) distribution of liquid assets

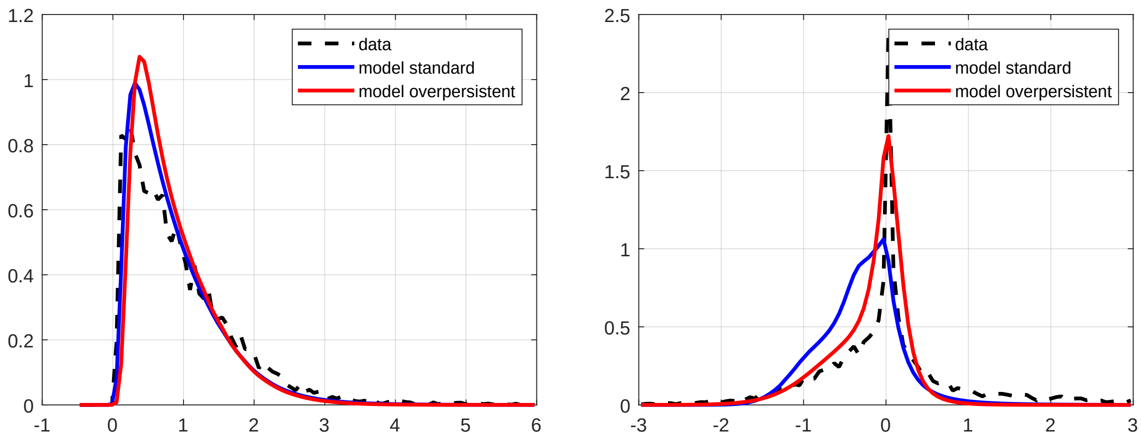

Note: The figure depicts the distribution for (a) durable goods and (b) liquid savings. Data distributions (dash-dotted black line) are compared to the distributions implied by model which allows for biased expectation (solid red line) and the model where expectations are assumed to be rational (solid blue line). The x-axis is normalised by the value of median quarterly income.

Finances and compare the credit card limit of an individual household to its quarterly income. On average in the period 1992-2010, households have a borrowing limit that is 56% of their quarterly income. We hence set $\kappa_y = 0.56$. Moreover, we further assume that households can borrow up to 80% against the value of their durable and set $\kappa_v = 0.8$. This is in line with Attanasio, Goldberg and Kyriazidou (2008) who report that the average finance share for households buying cars is 0.78.

To determine the depreciation rate δ and the proportional adjustment costs F^d we proceed as follows. The adjustment costs can be understood as the share of value a car loses just because it is sold to another person, i.e. the fraction of the purchase price which is not recovered if a car was resold immediately after the original purchase. We assume that this fraction is equal to 30% compared to the original value of the car and hence set $F^d = 0.3$. Furthermore, we assume that the resale value of a durable is negligible after 10 years. Given the adjustment costs F^d, this is the case for a quarterly depreciation rate of 5%. We therefore set $\delta = 0.05$.

Preference parameters The remaining five parameters are the preference parameters which affect the trade-off between non-durable consumption and the durable good $(\theta, \xi, \bar{d})$, risk aversion (γ) and the discount factor (β). The values of these parameters are chosen to match the aggregate distribution of liquid assets and the stock of durable goods in the data.

The data distributions we target have been obtained from the Survey of Consumer Finances (SCF), waves 1992-2010. The data counterpart for liquid assets is the sum of checking accounts, savings accounts, stocks, bonds, and mutual funds minus outstanding credit card debt after last payment and outstanding auto loans. Durable goods are defined as the current value of all vehicles belonging to the household. To eliminate effects of life-cycle savings we focus on the sample of vehicle owners aged 25-55.

The optimal parameter values are found using a grid search procedure. The resulting values are the discount factor $\beta = 0.9825$, risk aversion $\gamma = 1.5$, weight of durable goods $\theta = 0.075$, elasticity of substitution between durables and non-durables $\xi = 3$ and free durable services $\bar{d} = 0.5$.

Model Fit of Asset Distributions (targeted) Figure 7 shows that the model is able to replicate key features of the distributions of both durable goods and liquid assets. The model achieves a very good fit for the distribution of durable goods in the economy. In terms of liquid assets, the model succeeds in replicating the mass of households with zero liquid assets. It is important to stress that each of the two distributions is an infinite dimensional object and the model has only 5 parameters to achieve a good fit. The model struggles to replicate the thick right tail of the liquid assets distribution. In the model agents hold liquid assets for transactionary (due to the adjustment costs in durables) and precautionary reasons. It does not, however, capture life cycle motives for savings, nor does it include heterogeneity in preferences or heterogeneity in returns that households earn on their investments. Life-Cycle savings motives have been shown to help generate wealth inequality (see, e.g., De Nardi and Fella (2017) for a survey). Moreover, recent evidence shows that empirically, heterogeneity in returns is pronounced and can explain the large concentration of wealth at the top (see Fagereng et al. (2016), Bach, Calvet and Sodini (2017)). Hubmer, Krusell and Smith (2017) show that Bewley-type models like the one in this paper are not able to match the asset concentration at the top without adding heterogeneity in both preferences and returns. They also find that even with both of these sources of heterogeneity the models are unable to match the wealth holdings at the very top. Since our focus here is not on the top end of the wealth distribution we choose to abstract from these additional complexities.

Model Fit of Marginal Propensities to Consume (untargeted) Next we scrutinize the fit of the model by comparing the simulated MPCs with their empirical counterparts from the literature. MPCs were not targeted when determining the preference parameters so that this comparison can serve as a test for the overall fit of the model. Because the durable good is calibrated to represent cars (not housing), the model generates a wide set of predictions that can be brought to the data. We report separately the marginal propensity to consume in nondurable consumption (MPC), the marginal propensity to spend on durables (MPD) and the marginal propensity for total expenditures (MPE), where total expenditures combine nondurable consumption and durable expenditures but exclude adjustment costs.

The most relevant empirical estimates can be found in Johnson, Parker and Souleles (2006), Parker et al. (2013) and Misra and Surico (2014) for the reactions to the 2001 and 2008 stimulus payments in the U.S., as well as in Fagereng, Holm and Natvik (2018) who report the MPE from an ideal natural experiment of lottery winnings in Norway. The simulated marginal propensities to spend on the different goods have been constructed as the change in expenditures in reaction to a one-time, unanticipated transfer of 5% of median income. This size is comparable to the actual transfer people received in 2001 and 2008 in the U.S.. The technical details of the procedure to construct the MPC in the simulation are described in the online appendix.

Figure 8 shows the cumulative response in nondurable consumption (panel (a)) and total expenditures (panel (b)). The model predicts an average MPC on impact in nondurable consumption of just below 10% and that 38% of the transfer are on average spent on durables, resulting in a MPE of 47% of the transfer being spent on impact. Over the course of the first year almost 75% of the transfer are spent in total (25% is spent on nondurables). Over longer horizons, the contribution of nondurables increases.

How do these numbers compare to empirical studies?[10] In the quarter of the transfer, Souleles (1999, table 5) finds a MPE of 0.34 or 0.64 (depending on the estimation method). For the same horizon Parker et al. (2013, table 2) report an MPE of 0.516. The simulated MPE in the model of 0.47 lies well within one standard error of any of these estimates. At the 6 months horizon Fagereng, Holm and Natvik (2018, table 4) report that 52% of the lottery winnings are spent, which is very close to the simulated counterpart of 60%.

How about the split between nondurables and durables? Parker et al. (2013, table 2) report MPCs of 0.079 and 0.121 for "*strictly nondurables*" and "*nondurable*" spending, respectively. Similarly, Souleles (1999) finds the MPC to be 0.045-0.093 on strictly nondurables. The model matches these estimates with a MPC of just below 10% for nondurables. Turning to durable expenditures, Souleles (1999) report an MPD of 0.294-0.537 for all durables (0.166-0.24 for vehicles). Parker et al. (2013, panel E, table 7) find an MPD of 0.527 for vehicles. In contrast, Fagereng, Holm and Natvik (2018) find only 3% of lottery winnings is spent on cars and boats. The simulated MPD in the model of 0.38 is again well within the range of these empirical estimates.

[10] To compare our results with empirical studies we use point estimates reported in the literature. Needless to say that there is often large uncertainty around these estimates. For a recent overview of the empirical findings, see Carroll et al. (2017, table 1).

Figure 8: Reaction to an unexpected transfer

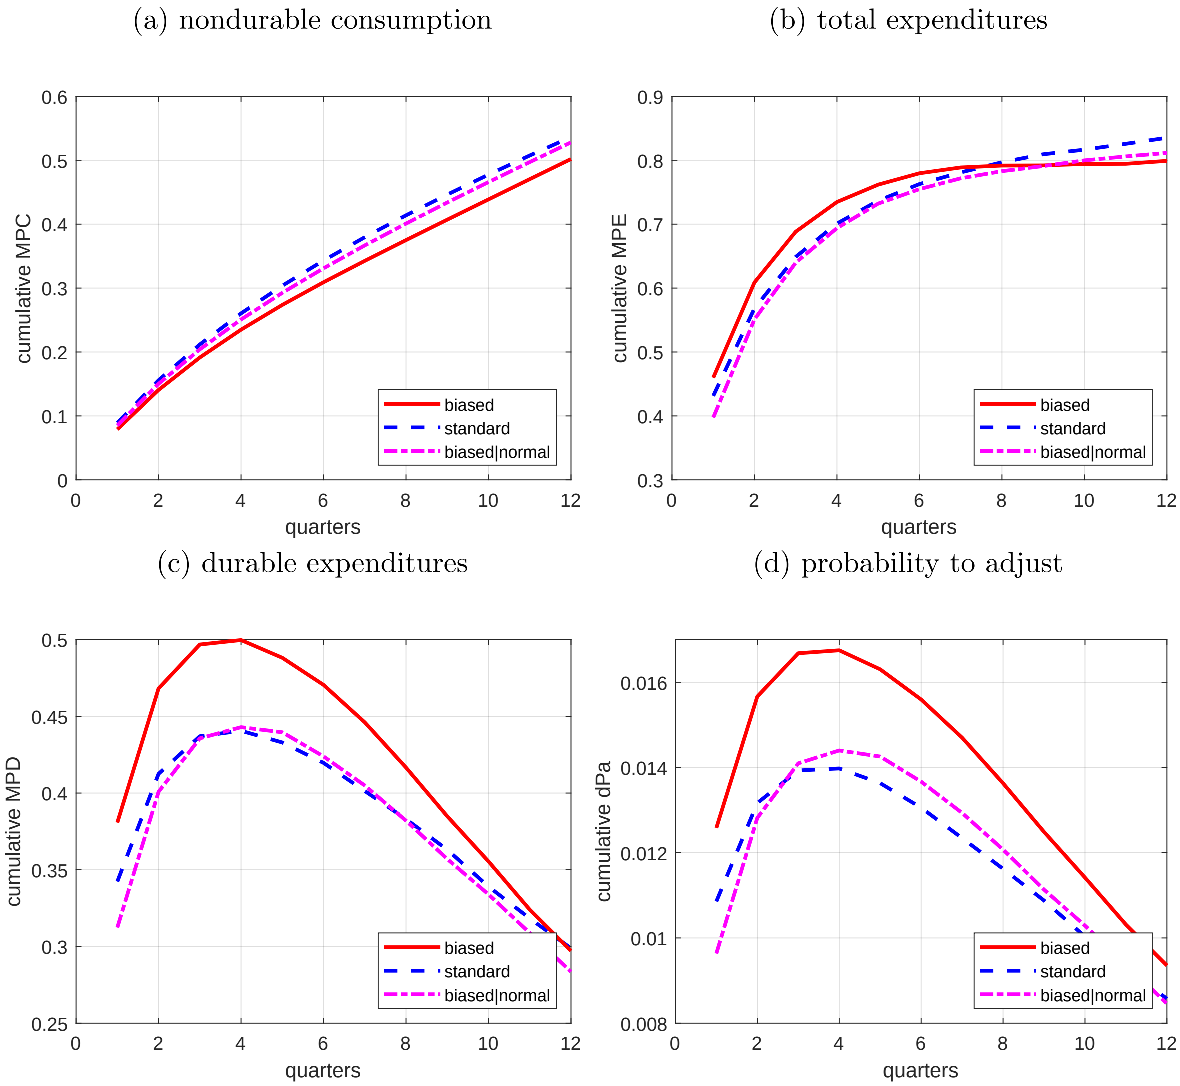

Note: The figure depicts how much of an unanticipated transfer households spend over time. Panel (a) depicts expenditures in nondurable consumption, panel (b) depicts total expenditure, panel (c) depicts durable expenditures, and panel (d) depicts the change in the probability to adjust the durable stock. All results are cumulative. The transfer size is equal to 5% of median quarterly household income.

Moreover, consistent with the empirical findings reported by Misra and Surico (2014), the model generates large heterogeneity in responses: First, the vast majority of households (over 95%) do not adjust their durables, neither with nor without the transfer, which results in their MPD being 0. Second, there is a small group of households (just below 3%) who were adjusting even in the absence of the transfer. For these households, the average MPD is roughly 1/3. Finally, there is an even smaller group of households (just above 1%) who would not have adjusted their durable stock without the transfer, but decide to do so when they receive the transfer. The transfer thus makes them move the adjustment date forward. For these households, the MPD is much larger than 1, because the size of the durable purchase is an order of magnitude larger than the size of the transfer. Cumulatively, the effect on durable expenditures peaks at one year, suggesting that the households who were induced to buy a car when they got the transfer were close to buying one even in the absence of the transfer. This shows that both the intensive and the extensive margin of durable purchases are operating and important for investigating MPDs.

To summarize, we conclude that the model with biased income expectations not only captures well the targeted distributions of liquid assets and durable goods but is also able to capture well the untargeted patterns of MPCs documented in the literature. This is true for both the overall expenditures as well as for the split between expenditures on durable and nondurable goods.

III.C Effects of Biased Income Expectations

In this section we show how the beliefs about income expectations affect the behavior of households in different income groups. To do so we compare the implications of the calibrated model with biased expectations to the implications of the same model under rational expectations (i.e. same parametrization). This highlights the effect of biased income expectations holding everything else equal.

We demonstrate that the overpersistence bias in expectations allows the model to fit the joint distribution of income and liquid assets. Above all, incorporating biased expectations reduces the amount low income households borrow, which is consistent with the data. Furthermore, we show how biased income expectations affect the consumption responses to unanticipated transfers. The overpersistence bias affects the MPC, MPD and MPE differentially across the income distribution. Low income households turn out to have a lower MPC in nondurables if they have biased expectations while the corresponding MPC of high income households is hardly affected by the beliefs. The differences in MPCs across the income distribution are hence smaller than what would be predicted under rational expectations. Moreover, the overpersistence bias makes the extensive margin of durable purchases more responsive to the transfer. This increases the MPE of households with biased expectations relative to that of rational agents for all but the very income-poorest households.

III.C.1 Effects on Behavior Across Different Income Groups

Figure 9 shows the distribution of durable goods and liquid assets for households in the lowest and highest income quintiles. The model is able to match the cross-sectional variation in durable holdings (panels (a) and (b)). This is true for both the model that allows for the

Figure 9: Distribution of assets across income groups

(a) durables, lowest quintile

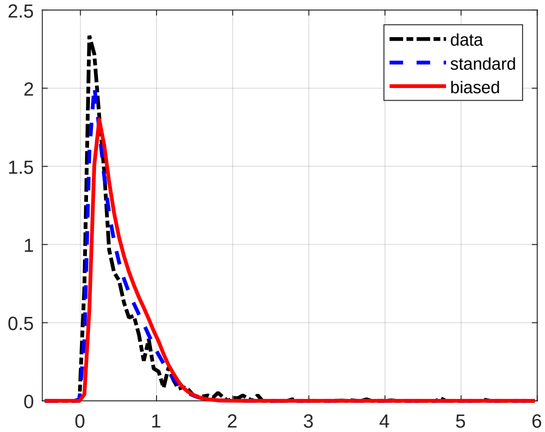

(b) durables, highest quintile

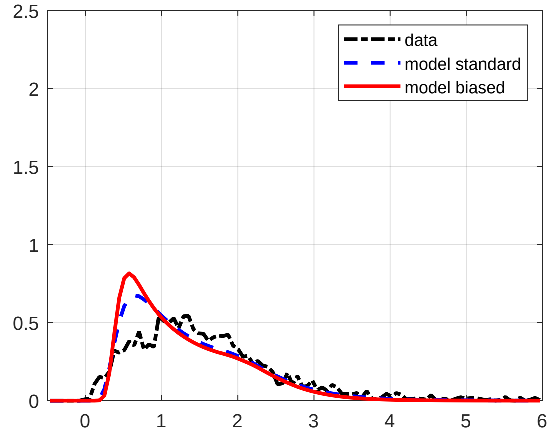

(c) liquid assets, lowest quintile

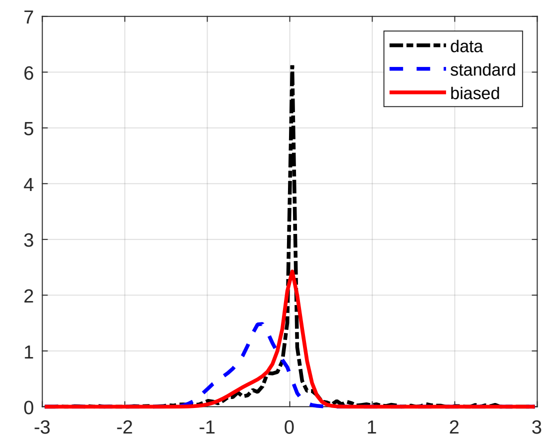

(d) liquid assets, highest quintile

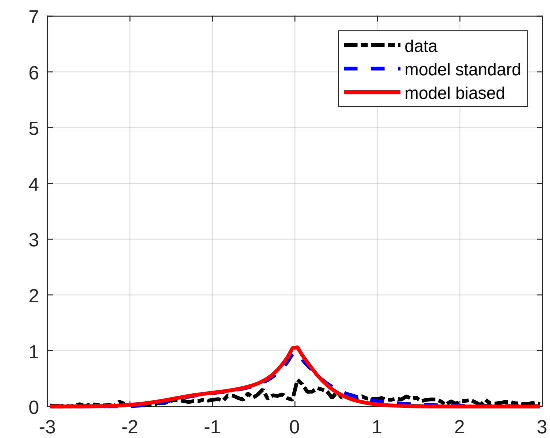

Note: The figure depicts the distribution of durable goods and liquid assets for different income quintiles. The panels show the data distribution (dash-dotted black line) against the model distribution when households have biased expectations (solid red line). For comparison, the distribution under rational expectations is also plotted (dashed blue line).

expectation bias and for the fully rational model. In terms of durable holdings, biased expectations hence do not change the distributions much compared to the distributions implied by rational expectations.

However, this is not true for the distribution of liquid assets. Figure 9, panels (c) and (d) shows the distribution for liquid assets for the two different income quintiles. While the distribution in the highest income group is not much affected by biased income expectations, the behavior of the lowest income group depends on what households believe about their future income. Low income households with biased beliefs are too pessimistic about their future income. They are therefore less willing to borrow even though their borrowing constraint is not binding. Figure 9(c) shows that this mechanism allows the model with biased income expectations to fit the empirical distribution of liquid assets in the lowest income group very well. It is important to note that with biased beliefs, low income households choose not to borrow more even though they could. If people had rational expectations instead, the model would predict counterfactually large amounts of borrowing (mode of -0.5 versus 0 in the data).

Figure 10 shows in more detail the heterogeneity of how the overpersistence bias alters the behavior. Panel (a) displays the probability of adjusting the durable stock for each percentile of the income distribution. It shows that the likelihood of adjusting the durable stock increases monotonically as income increases. However, allowing for biased expectations flattens this income gradient: Households in the bottom 80% of the income distribution adjust their durable more frequently than their rational expectation counterparts. Panels (b) to (d) detail the behavior given adjustment. Allowing for biased expectations very slightly reduces the size of the purchased durables for the whole income distribution except for the very lowest income households. More striking, however, are the differences in consumption and savings at time of purchase: Biased expectations reduce the average amount of borrowing for lower income households and increase the amount of borrowing for high income households. At the same time, lower income households consume less while high income households consume more than their rational expectation counterparts. This effect on consumption and borrowing is the direct effect of biased expectations: Low income households are too pessimistic about their future income and therefore want to save more (or borrow less). High income households, in contrast, are too optimistic and are hence willing to spend more and borrow to finance these expenditures.

III.C.2 Implications for Marginal Propensity to Consume

Turning to the effects of the bias on MPCs, figure 8, panel (a) shows that the average MPC in nondurable goods is lower for households with overpersistence bias compared to rational households. At the same time, however, durable expenditures are more responsive to transfers if households have biased beliefs (figure 8, panel (c)). The reason is that under biased beliefs more households are induced to move their adjustment date of durable stocks forward (panel (d)) while they spent less of the transfer conditional on adjustment. They thus react more on the extensive margin while reacting less on the intensive margin. Note that the differences in behavior are the result of two sources. First, biased beliefs imply that households have different expectations about their future income compared to rational agents. Second, they already had biased expectations in the past and hence made different

Figure 10: Income gradient of probability to adjust and behavior at the time of adjustment

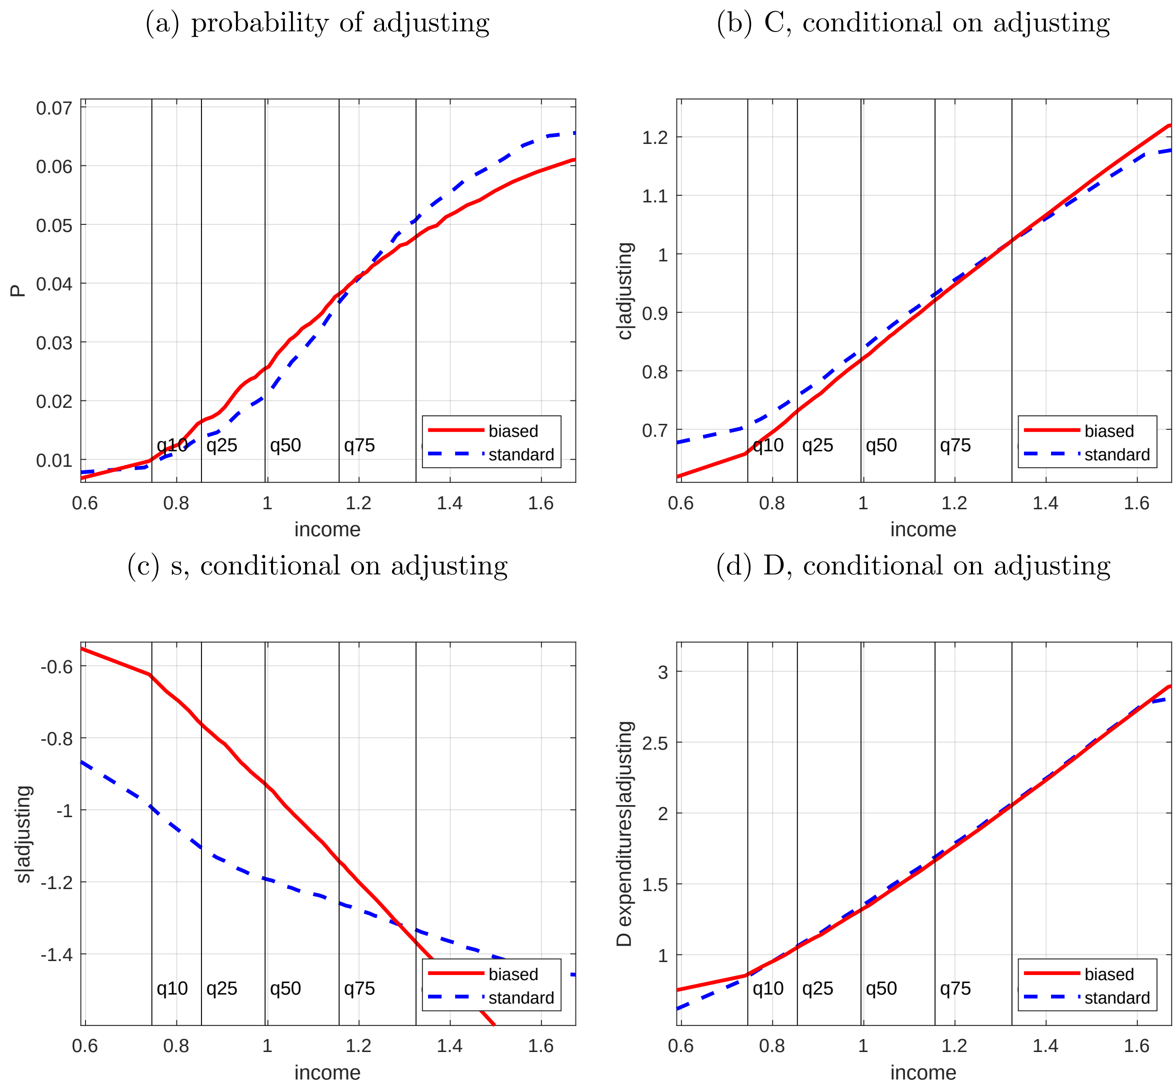

Note: The figure depicts the average behaviour across the income distribution under different expectation scenarios: the red line depicts the behaviour under biased expectations, the dashed blue dashed line depicts the behaviour under rational expectations and the magenta dash-dotted line shows what the behaviour of the overpersistent population would be if they were given the liquid assets and durable stock of the rational agents. The figure displays the behaviour within the 1st and 99th percentile of the income distribution, and the vertical lines denote the 10th, 25th, 50th, 75th and 90th percentile, respectively.

consumption-savings decisions. They therefore have a different asset position than their rational expectations counterparts. Figure 8 shows that most of the differences in the durable consumption response are due to the second effect. The dash-dotted magenta line depicts the response in durable expenditures if households have biased beliefs going forward but currently have the asset position of rational agents. We see that this response is much closer to the response in the rational model which indicates that biased beliefs mainly alter durable consumption responses through their effect on the asset position.

How do biased expectations affect the distribution of MPCs? Figure 11 depicts the differential effect of the bias on the MPC in nondurables for low and high income households. Panel (a) shows that low income households with biased expectations have an MPC in nondurables that is between 3-11 percentage points lower than the MPC of rational households, depending on the horizon. These differences are the result of the same two forces as explained above: First, low income households with biased expectations are too pessimistic about their income going forward. This implies that they are more cautious in spending the transfer payment and more likely to save out of it. Second, they have a different asset position compared to their rational expectations counterparts. Since they have already been too pessimistic in the past they are less likely to be close to the borrowing constraint. Figure 11 shows that both of these forces contribute to the reduction in the MPC of biased households and that the magnitude of the two effects is similar. Households in the highest income quintile, on the other hand, spend about the same fraction of the transfer payment on nondurables whether they have biased expectations or not (panel (b)). Figure 11, panels (c) and (d) display the corresponding impulse response functions for total expenditures. In this case, low income households have a similar MPE during the first 1.5 years after the transfer whether they have biased expectations or not. This is because the effect of differing expectations and differing asset position cancel each other out. For high income households, on the other hand, both effects are positive. The overpersistence bias hence increases the MPE for the high income group.

Figure 12 shows in more detail the heterogeneous effects of biased expectations across the income distribution. It depicts the MPCs, MPEs, MPDs on impact, and the change in propensity to adjust due to an unexpected transfer. Panel (a) shows that under both rational and biased expectations the MPC in nondurable consumption is falling in income. Under biased expectations this income gradient is similar across the whole income distribution. Under rational expectation, in contrast, the MPC is much larger for the lowest income households. This emphasizes that the bias has the most effect on the lowest income households. Panel (b) shows that the MPE is increasing with income across the whole income distribution. This increasing slope is driven by durable expenditures (panel (c)). Panel (d) further shows that as a reaction to the transfer the likelihood of adjustment increases, in particular for the middle income groups. All else equal, having biased income expectations increases the effect that the transfer has on the likelihood to adjust for the bottom 75% of the income distribution. This translates into a higher MPD for these income groups. The average MPE is also higher for these income groups except for the lowest income group. For these households, the decrease in MPC is strong enough to overcompensate the increase in

Figure 11: Cumulative MPC and MPE out of unexpected transfer by income

(a) MPC, lowest income quintile (b) MPC, highest income quintile

(c) MPE, lowest income quintile (d) MPE, highest income quintile

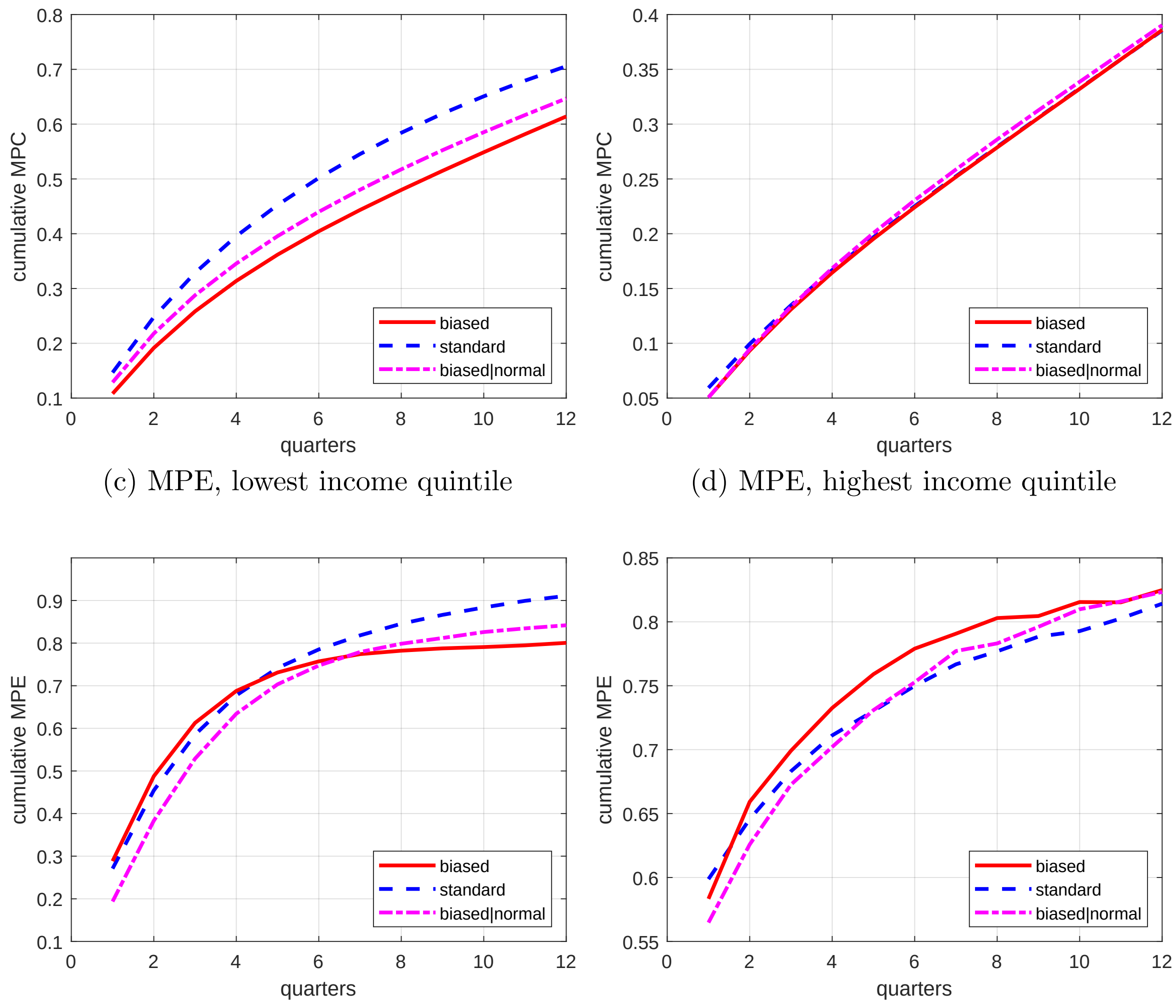

Note: The figure depicts the fraction of an unanticipated one-time transfer payment that is spent on non-durable consumption and total expenditures under different expectation scenarios: the red line depicts the behaviour under biased expectations, the dashed blue dashed line depicts the behaviour under rational expectations and the magenta dash-dotted line shows what the behaviour of the overpersistent population would be if they were given the liquid assets and durable stock of the rational agents. Panels (a) and (c) show the results for the lowest income quintile, panels (b) and (d) for the highest income quintile. The transfer size is equal to 5% of median quarterly income in the economy.

MPD. For the lowest income group, even the MPE is thus lower under biased expectations than under rational expectations.

III.D Aggregate Implications for Government Stimulus Programs

In the previous section we have analysed the effects of biased expectations while holding all else equal, including preference parameters and the market environment. We have demon-

Figure 12: Income gradient of marginal propensities to consume

(a) nondurable consumption (b) total expenditures

(c) durable expenditures (d) probability of adjustment

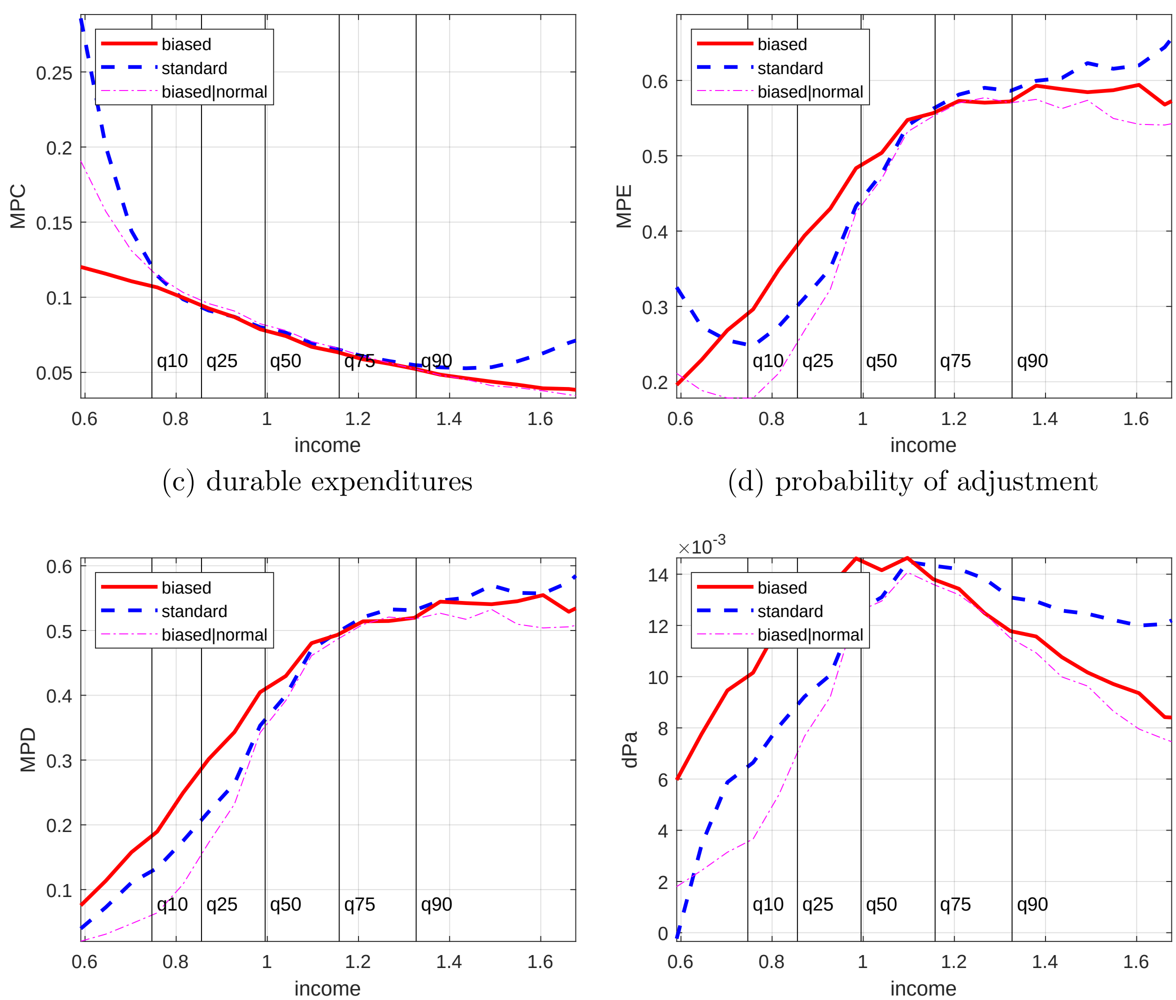

Note: The figure depicts the average MPC across the income distribution under different expectation scenarios: the red line depicts the MPC under biased expectations, the dashed blue dashed line depicts the MPC under rational expectations and the magenta dash-dotted line shows what the MPC of the overpersistent population would be if they were given the liquid assets and durable stock of the rational agents. The figure displays the behavior within the 1st and 99th percentile of the income distribution, and the vertical lines denote the 10th, 25th, 50th, 75th and 90th percentile, respectively.

strated that the calibrated model with biased expectations can generate the limited borrowing behaviour of low income households. However, in the fully rational counterpart to this model, these households borrow too much compared to the data. In order to reduce this excessive borrowing one can make the borrowing behaviour an explicit target for the calibration and calibrate an exogenous borrowing constraint to fit this data moment. In this section we follow this approach for both the model with biased and the model with rational expectations. We show that a model with biased income expectations requires less restrictive borrowing constraints than the rational model to fit the data. This has important implications for MPCs and hence for the effectiveness of fiscal stimulus programs. These government policies are a popular instrument during recessions to boost household consumption in order to stabilize the overall economy. In both recent recessions in 2001 and 2008, the U.S. government employed this strategy by giving households one-off cash transfers. We find that the rational model predicts such stimulus policies to be substantially more effective than what the model with biased income expectations predicts.

We make the following change to the model setup: We replace the endogenous borrowing limit $\kappa_y P_t + \kappa_v d_t$, which depended on a household's income and durable stock with a constant $\underline{s}$ that is the same for all households. Equation (15) hence becomes

$$R(s_t) = [1 + r(s_t)]s_t \text{ where } r(s_t) = \begin{cases} r^l & \text{if } s_t > 0 \\ r^b & \text{if } \underline{s} \leq s_t \leq 0 \end{cases} \tag{18}$$

We leave all other parameters the same as in the main calibration and determine the level of $\underline{s}$ that is needed to fit the empirical share of households with positive liquid assets from the Survey of Consumer Finances (as before including the outstanding credit card balance and auto loans). Figure 13 displays this share for varying levels of the borrowing limit for both the rational model and the model with biased income expectations. While in the biased model the required borrowing limit is –0.3675, the rational model needs a much tighter borrowing constraint of –0.145 to fit this data moment. With a median income of one, this difference corresponds to about 1/4 of median quarterly income.

This difference in borrowing limits amplifies the effects of expectations on the level and distribution of MPCs. Figure 14 shows that for all possible levels of the borrowing constraint, the rational model results in a higher aggregate MPC than the biased model. Moreover, the aggregate MPC increases as borrowing constraints are tightened. Taken together both effects imply that the rational model results in a aggregate MPC that is 50% higher compared to the MPC in the model with biased expectations: it is 0.18 in the rational model while the model with biased expectations leads to an MPC in nondurables of 0.12. To put this into perspective, the "Economic Stimulus Act" of 2008 provided about 100 billion US dollars in tax relief to households (Shapiro and Slemrod, 2009; Parker et al., 2013). In this context, the rational model would predict that this stimulus package increases nondurable consumption by 18 billion US dollars. The model with biased expectations, on the other hand, only predicts and increase of 12 billion US dollars.

Moreover, the distribution of MPCs across the income distribution is affected by the different borrowing constraints. Figure 14, panel (b), displays the ratio between the average

Figure 13: Share of households with non-negative liquid assets by borrowing constraint

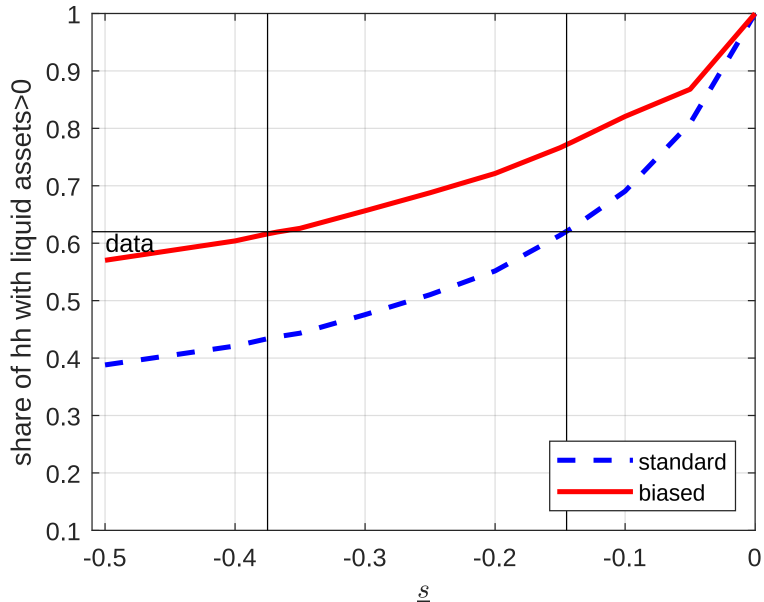

Note: The figure depicts the share of households with non-negative liquid assets for model specifications with varying borrowing constraints $\underline{s}$. The solid red line refers to the model with biased income expectations; the dashed blue line refers to the rational model. The horizontal line depicts the empirical value computed from the Survey of Consumer Finances (at 0.62), the vertical lines (at $\underline{s} = -0.3675$ and $\underline{s} = -0.145$ mark the borrowing constraint required in either model to match the data.

MPC of households in the lowest income quintile and the average MPC in the highest income quintile. The biased model predicts this relative MPC to be 2.2. This value is well within the range of empirical estimates: Johnson, Parker and Souleles (2006) and Parker et al. (2013) obtain point estimates of this MPC ratio of 2.33-2.99 and 1.16 for the stimulus payments in the United States in 2001 and 2008, respectively.[11] In contrast, the rational model predicts that low income households spend almost 6 times as much as high income households out of the transfer. In the previous section we had shown that all else equal the bias decreases the MPC in nondurables for the lowest income households. The large relative MPC here demonstrates that the tighter required borrowing constraint in the rational model further amplifies this effect on the MPC of low income households.

Figure 14, panels (c) and (d) display the corresponding results for total expenditures. While the level of the aggregate MPE is comparable in the rational and biased models (0.37 vs 0.34), the distribution of MPEs remains very different in the two models. The rational model predicts that the MPE of low income households is about the same as that of high income households (i.e. a relative MPE of 1). In contrast, in the model with biased income

[11] Johnson, Parker and Souleles (2006) define income groups as: low < \$34K, high > \$69K. Parker et al. (2013) define income groups as: low < \$32K, high > \$75K.

Figure 14: MPC and MPE for different levels of the borrowing constraint

(a) Aggregate MPC

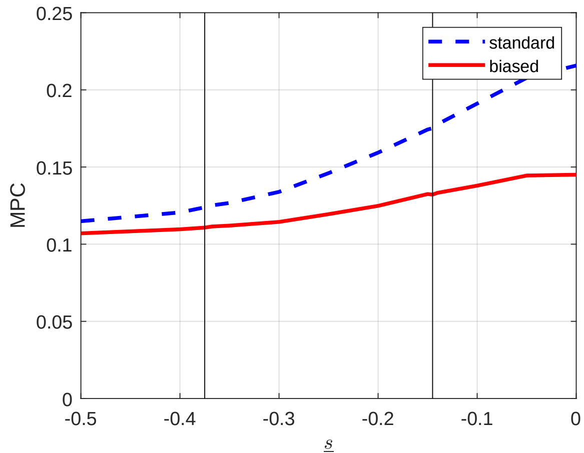

(b) Relative MPC

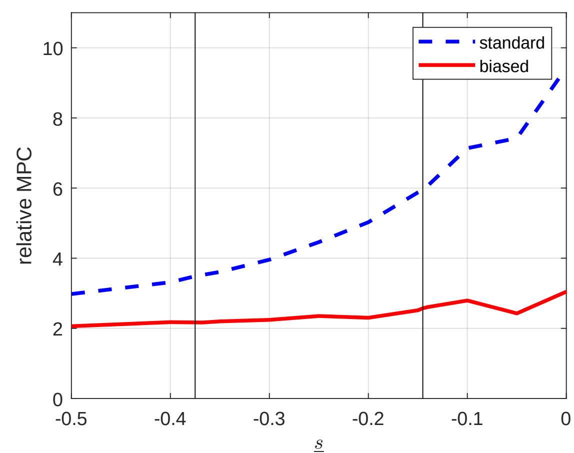

(c) Aggregate MPE

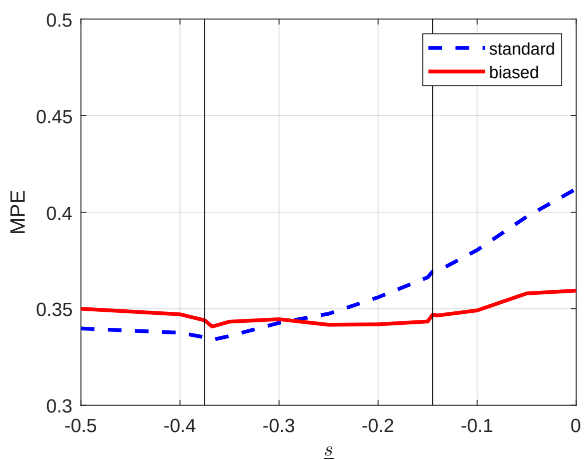

(d) Relative MPE

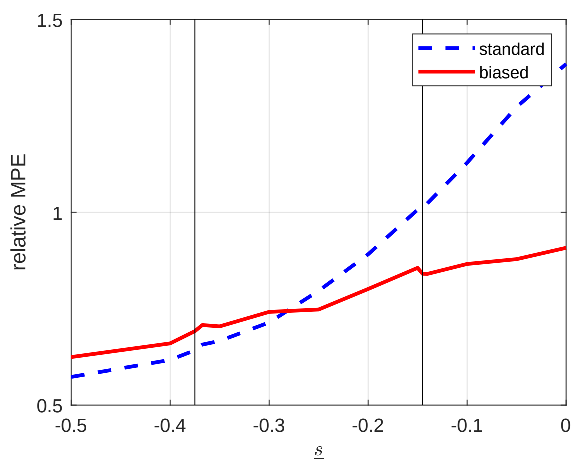

Note: The figure depicts the aggregate and relative MPC and MPE for model specifications with varying borrowing constraints $\underline{s}$. Relative MPC (relative MPE) are the ratio of the MPC (MPE) of the lowest income quintile over that of the highest income quintile. The solid red line refers to the model with biased income expectations; the dashed blue line refers to the rational model. The vertical lines $\underline{s} = -0.3675$ and $\underline{s} = -0.145$ mark the borrowing limit required to match the data by the rational and biased model, respectively.

expectations low income households spend only 70% of what high income households spend out of the transfer.

The difference in MPC (and MPE) between high and low income households is important for the effectiveness of government stimulus policies. Stimulus payments have to be financed in some way, which is often done through taxes. Since high income households typically pay higher taxes than low income households, stimulus payments are a form of redistribution. How much aggregate consumption increases due to this transfer therefore depends on the

Figure 15: Aggregate effects of a redistributive policy

(a) Nondurable consumption (b) Total expenditures

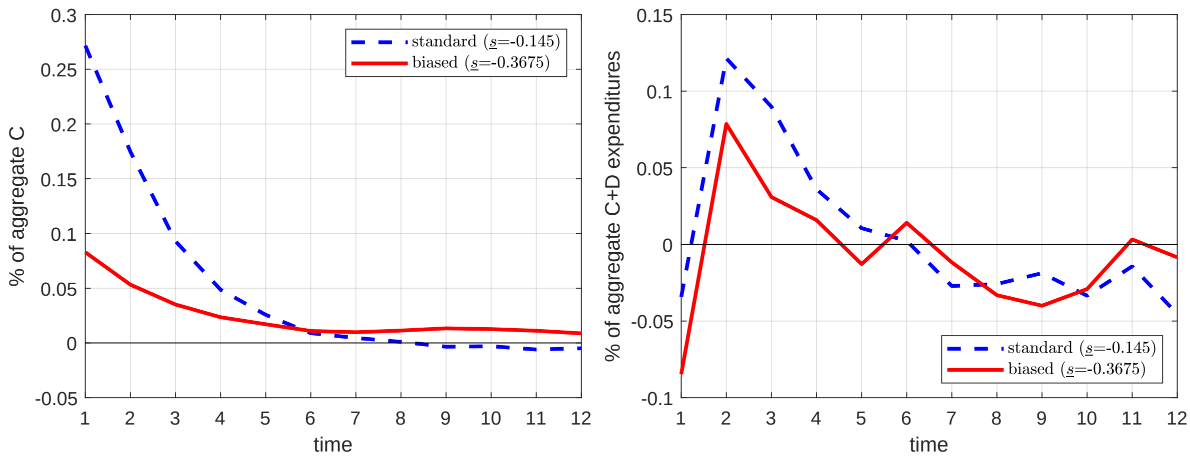

Note: The figure depicts the aggregate impulse-response function of nondurable consumption and total expenditures to a redistributive policy: households below median income receive a transfer while households above median income pay for the transfer. The figure shows the results for the biased model (solid red line) and rational model (dashed blue line). Magnitudes are expressed as percentage increase in nondurable consumption or total expenditures, respectively, relative to the level without the policy.

ratio between the MPC (and MPE) of low income households and high income households.

To illustrate the importance of the distribution of marginal propensities across income, we compute the reactions to a theoretical redistributive stimulus policy with a balanced budget. We assume that all households with income below the median income receive a transfer (or tax rebate) of 5% of median household income. This transfer is financed by levying a lump sum tax of the same magnitude on all households above median income. Figure 15 shows the resulting aggregate impulse response functions. The model with overpersistence bias predicts that this policy increases nondurable consumption on impact by 8%. In contrast, according to the rational model we would expect this policy to increase nondurable consumption by 27%, that is almost 3.5 times the effect in the biased model. In terms of total expenditures we see an almost equally large discrepancy between the two models: On impact both models predict a fall in total expenditures. This fall is more pronounced in the model with biased expectations than in the rational model (–8% vs –3%). In the subsequent quarter the model with biased expectations predicts total expenditures to increase by 7% while the rational model expects an increase in total expenditures of 12%. The rational model thus predicts this policy to be more effective than what the model with biased expectations predicts. This is true both for nondurable consumption and for total expenditures.

To summarize, we find that the rational model requires a tighter borrowing constraint than the biased model to fit the data. This tighter borrowing limit amplifies the differences between the biased and rational models. The rational model generates a larger MPC and MPE for low income households relative to high income households than the model that accounts for the overpersistence bias. The rational model therefore predicts stimulus payments to be more effective than what the model with biased expectations predicts.

IV Conclusion

In this paper we investigate the role of income expectations on consumption behavior of households. We document a systematic bias in income expectation, show how it can be formally incorporated into the process of expectation formation and investigate its implications for consumption-saving decisions in a quantitative model.

Using household level data from the Michigan Surveys of Consumers, we find that households with high income today tend to overestimate their future income and those with low income underestimate their future income. We argue that this feature of expectation bias can be explained by households overestimating the persistence of their income process. This overpersistence belief is consistent with the observation that people fail to sufficiently appreciate regression to mean. This observation is not new to behavioral economics and psychology (see Kahneman and Tversky (1973) and Kahneman (2012, chapter 17)). However, to the best of our knowledge this paper is the first to quantify the extent of the bias in income expectations and investigate its implications for consumption decisions using a quantitative model.

We proceed by exploring an economy where households exhibit the same expectation biases as we observe in the data. The model we build is an otherwise standard partial equilibrium consumption model with a durable asset subject to adjustment costs. Income expectation biases of the magnitude seen in the data significantly affect the distribution of liquid assets in the cross section. Low income households with biased beliefs are too pessimistic about their future income and are hence unwilling to borrow to smooth consumption. In contrast, households with high income turn out to have similar portfolios of durable goods and liquid savings whether they have biased income expectations or not. This prediction of the model with biased beliefs is in line with the distribution of liquid assets in the data.

The paper further showed that if instead we calibrate the exogenous borrowing constraint in a rational model, the rational model needs a tighter borrowing limit to fit the data than the model with biased expectations. This tighter borrowing limit amplifies the effects of biased beliefs. In particular, the model with biased expectations leads to a substantially lower relative MPC of low income households to that of high income households. If stimulus payments are financed through taxes (which are predominantly paid by high income households), stimulus payments are a form of redistribution. How effective these programs are hence depends on the distribution of MPCs across the income distribution. The paper showed that the differences in the distribution of MPCs translate into economically meaningful differences in the assessment of fiscal stimulus policies: According to the biased model stimulus payments are substantially less effective than according to the rational model.

We believe that our empirical finding opens an avenue for further research in two main areas. First, while the available data from the Michigan Surveys of Consumers allows us to document patterns in income expectation biases, the data set has an important limitation: it has only a very short panel dimension. This limitation makes it impossible to follow the same households and their expectations over time. Using the Michigan Surveys of Consumers we are therefore unable to investigate in detail the process of expectation formation and expectation updating. Other existing panel surveys do not include enough information to analyze expectation biases in individual income expectations. To learn more about how income expectations are formed it thus seems very important to collect new data both on income expectations and on the corresponding realizations in a panel survey.

Second, our analysis shows that there are substantial movements in income expectation errors at the business cycle frequency. This suggests a role for income expectation errors for macroeconomic business cycle analysis. In the present paper we have focused on the cross-sectional patterns of expectation errors. In future work it would be interesting to study these business cycle movements in expectation errors and analyze the effects that household income expectations have for the amplification of other types of macroeconomic shocks.

References

Andolfatto, David, Scott Hendry, and Kevin Moran. 2008. "Are inflation expectations rational?" *Journal of Monetary Economics*, 55(2): 406 – 422.

Andrade, Philippe, and Hervé Le Bihan. 2013. "Inattentive professional forecasters." *Journal of Monetary Economics*, 60(8): 967 – 982.

Armantier, Olivier, Wändi Bruine de Bruin, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar. 2015. "Inflation Expectations and Behavior: Do Survey Respondents act on their Beliefs?" *International Economic Review*, 56(2): 505–536.

Armona, Luis, Andreas Fuster, and Basit Zafar. 2018. "Home Price Expectations and Behaviour: Evidence from a Randomized Information Experiment." *The Review of Economic Studies*, 86(4): 1371–1410.

Attanasio, Orazio P., Pinelopi Koujianou Goldberg, and Ekaterini Kyriazidou. 2008. "Credit Constraints in the Market for Consumer Durables: Evidence from Micro Data on Car Loans." *International Economic Review*, 49(2): 401–436.

Bach, Laurent, Laurent E. Calvet, and Paolo Sodini. 2017. "Rich Pickings? Risk, Return, and Skill in the PPortfolio of the Wealthy."

Bachmann, Rüdiger, Tim O. Berg, and Eric R. Sims. 2015. "Inflation Expectations and Readiness to Spend: Cross-Sectional Evidence." *American Economic Journal: Economic Policy*, 7(1): 1–35.

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer. 2015. "X-CAPM: An extrapolative capital asset pricing model." *Journal of Financial Economics*, 115(1): 1 – 24.

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer. 2018. "Extrapolation and bubbles." *Journal of Financial Economics*, 129(2): 203 – 227.

Barnard, John, and Donald B. Rubin. 1999. "Small-sample degrees of freedom with multiple imputation." *Biometrica*, 86(4): 948–955.

Berger, David, and Joseph Vavra. 2015. "Consumption Dynamics During Recessions." *Econometrica*, 83(1): 101–154.

Bewley, Truman F. 1986. "Stationary monetary equilibrium with a continuum of independently fluctuating consumers." In *Contributions to Mathematical Economics in Honor of Gerard Debreu.* , ed. Werner Hildebrand and Andreu Mas-Colell, 79–102. North-Holland.

Bhandari, Anmol, Jaroslav Borovicka, and Paul Ho. 2019. "Survey Data and Subjective Beliefs in Business Cycle Models."

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2018. "Diagnistic Expectations and Credit Cycles." *Journal of Finance*, 73(1): 199–227.

Broer, Tobias, and Alexandre Kohlhas. 2018. "Forecaster (Mis-)Behavior." *CEPR Discussion Paper DP12898.*

Carroll, Christopher D. 2003. "Macroeconomic Expectations of Households and Professional Forecasters." *The Quarterly Journal of Economics*, 118(1): 269–298.

Carroll, Christopher D., Edmund Crawley, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White. 2018. "Sticky Expectations and Consumption Dynamics."

Carroll, Christopher, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White. 2017. "The Distribution of Wealth and the Marginal Propensity to Consume." *Quantitative Economics*, 8(3): 977–1020.

Case, Karl E., Robert J. Shiller, Anne K. Thompson, David Laibson, and Paul Willen. 2012. "What Have They Been Thinking? Homebuyer Behavior in Hot and Cold Markets [with Comments and Discussion]." *Brookings Papers on Economic Activity*, 265–315.

Cavallo, Alberto, Guillermo Cruces, and Ricardo Perez-Truglia. 2017. "Inflation Expectations, Learning, and Supermarket Prices: Evidence from Survey Evidence." *American Economic Journal: Macroeconomics*, 9(3): 1–35.

Coibion, Olivier, and Yuriy Gorodnichenko. 2012. "What Can Survey Forecasts Tell Us about Information Rigidities?" *Journal of Political Economy*, 120(1): 116–159.

Coibion, Olivier, and Yuriy Gorodnichenko. 2015. "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts." *American Economic Review*, 105(8): 2644–78.

Coibion, Olivier, Yuriy Gorodnichenko, and Saten Kumar. 2015. "How Do Firms Form Their Expectations? New Survey Evidence."

Cutler, David M., James M. Poterba, and Lawrence H. Summers. 1990. "Speculative Dynamics and the Role of Feedback Traders." *American Economic Review Papers & Proceedings*, 80(2): 63–68.

Das, Marcel, and Arthur van Soest. 1999. "A panel data model for subjective information on household income growth." *Journal of Economic Behavior & Organization*, 40(4): 409 – 426.

Das, Sreyoshi, Camelia M. Kuhnen, and Stefan Nagel. 2017. "Socioeconomic Status and Macroeconomic Expectations."

Deaton, Angus. 1991. "Saving and Liquidity Constraints." *Econometrica*, 59(5): 1221–1248.

De Bruin, Wändi Bruine, Charles F. Manski, Giorgio Topa, and Wilbert van der Klaauw. 2011. "Measuring consumer uncertainty about future inflation." *Journal of Applied Econometrics*, 26(3): 454–478.

DeLong, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Waldman Robert J. 1990. "Positive Feedback Investment Strategies and Destabilizing Rational Speculation." *The Journal of Finance*, 45(2): 379–395.

De Nardi, Mariacristina, and Giulio Fella. 2017. "Saving and wealth inequality." *Review of Economic Dynamics*, 26: 280–300.

D'Haultfoeuille, Xavier, Christophe Gaillac, and Arnaud Maurel. 2018. "Rationalizing Rational Expectations? Tests and Deviations."

Dominitz, Jeff. 1998. "Earnings Expectations, Revisions, and Realizations." *The Review of Economics and Statistics*, 80(3): 374–388.

Dominitz, Jeff, and Charles Manski. 1997. "Using Expectations Data to Study Subjective Income Expectations." *Journal of the American Statistical Association*, 92: 855–867.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri. 2016. "Heterogeneity and Persistence in Returns to Wealth."

Fagereng, Andreas, Martin B. Holm, and Gisle J. Natvik. 2018. "MPC Heterogeneity and Household Balance Sheets."

Fuhrer, Jeffrey. 2017. "Intrinsic Expectations Persistence."

Gennaioli, Nicola, and Andrei Shleifer. 2010. "What Comes to Mind." *Quartlery Journal of Economics*, , (125): 1399–1433.

Gennaioli, Nicola, Yueran Ma, and Andrei Shleifer. 2016. "Expectations and Investment." *NBER Macroeconomics Annual*, 30(1): 379–431.

Gerardi, Kristopher, Andreas Lehnert, Shane M. Sherlund, and Paul Willen. 2008. "Making Sense of the Subprime Crisis." *Brookings Papers on Economic Activity*, 2008: 69–145.

Greenwood, Robin, and Andrei Shleifer. 2014. "Expectations of Returns and Expected Returns." *The Review of Financial Studies*, 27(3): 714–746.

Harmenberg, Karl, and Erik Öberg. 2017. "Consumption Dynamics under Time-varying Unemployment Risk."

Hubmer, Joachim, Per Krusell, and Anthony Smith. 2017. "The Historical Evolution of the Wealth Distirbution: A Quantitative-Theoretical Investigation."

Jappelli, Tullio, and Luigi Pistaferri. 2014. "Fiscal Policy and MPC Heterogeneity." *American Economic Journal: Macroeconomics*, 6(4): 107–36.

Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles. 2006. "Household Expenditure and the Income Tax Rebates of 2001." *American Economic Review*, 96(5): 1589–1610.

Kahneman, Daniel. 2012. *Thinking, Fast and Slow.* Penguin.

Kahneman, Daniel, and Amos Tversky. 1973. "On the Psychology of Prediction." *Psychological Review*, 80(4): 237–251.

Kaplan, Greg, and Giovanni L. Violante. 2014. "A Model of the Consumption Response to Fiscal Stimulus Payments." *Econometrica*, 82(4): 1199–1239.

Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante. 2016. "Non-durable Consumption and Housing Net Worth in the Great Recession: Evidence from Easily Accessible Data."

Kohlhas, Alexandre, and Ansgar Walther. 2018. "Asymmetric Attention."

Kuchler, Theresa, and Basit Zafar. 2018. "Personal Experiences and Expectations about Aggregate Outcomes."

Mackowiak, Bartosz, and Mirko Wiederholt. 2009. "Optimal Sticky Prices under Rational Inattention." *American Economic Review*, 99(3): 769–803.

Malmendier, Ulrike, and Stefan Nagel. 2015. "Learning from Inflation Experiences." *The Quarterly Journal of Economics.*

Mankiw, N. Gregory, and Ricardo Reis. 2002. "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *The Quarterly Journal of Economics*, 117(4): 1295–1328.

Massenot, Baptiste, and Yuri Pettinicchi. 2018. "Can Households See into the Future? Survey Evidence from the Netherlands."

Mian, Atif, Kamalesh Rao, and Amir Sufi. 2013. "Household Balance Sheets, Consumption, and the Economic Slump." *Quarterly Journal of Economics*, 128(4): 1687–1726.

Misra, Kanishka, and Paolo Surico. 2014. "Consumption, Income Changes and Heterogeneity: Evidence from Two Fiscal Stimulus Programmes." *American Economic Journal: Macroeconomics*, 6: 84–106.

Oh, Hyunseung, and Ricardo Reis. 2012. "Targeted transfers and the fiscal response to the great recession." *Journal of Monetary Economics*, 59(S): S50–S64.

Parker, Jonathan A. 2017. "Why Don't Households Smooth Consumption? Evidence from a $25 Million Experiment." *American Economic Journal: Macroeconomics*, 9(4): 153–183.

Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson, and Robert McClelland. 2013. "Consumer Spending and the Economic Stimulus Payments of 2008." *American Economic Review*, 103(6): 2530–53.

Piazzesi, Monika, and Martin Schneider. 2009. "Momentum Traders in the Housing Market: Survey Evidence and a Search Model." *The American Economic Review*, 99(2): 406–411.

Piazzesi, Monika, Juliana Salomao, and Martin Schneider. 2015. "Trend and Cycle in Bond Premia."

Reiter, Jerome P. 2007. "Small-sample degrees of freedom for multi-component significance tests with multiple imputation for missing data." *Biometrica*, 94(2): 502–508.

Rubin, Donald. 1987. *Multiple Imputation for Nonresponse in Surveys.* Wiley.

Shapiro, Matthew D., and Joel Slemrod. 2009. "Did the 2008 Tax Rebates Stimulate Spending?" *American Economic Review*, 99(2): 374–79.

Sims, Christopher A. 2003. "Implications of rational inattention." *Journal of Monetary Economics*, 50(3): 665 – 690. Swiss National Bank/Study Center Gerzensee Conference on Monetary Policy under Incomplete Information.

Souleles, Nicholas S. 1999. "The Response of Household Consumption to Income Tax Refunds." *American Economic Review*, 89(4): 947–958.

Souleles, Nicholas S. 2004. "Expectations, Heterogeneous Forecast Errors, and Consumption: Micro Evidence from the Michigan Consumer Sentiment Surveys." *Journal of Money, Credit and Banking*, 36(1): 39–72.

Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron. 2004. "Cyclical Dynamics in Idiosyncratic Labor Market Risk." *Journal of Political Economy*, 112(3): 695–717.

Vellekoop, Nathanael, and Mirko Wiederholt. 2017. "Inflation Expectations and Choices of Households."

Wiswall, Matthew, and Basit Zafar. 2015. "Determinants of College Major Choice: Identification using an Information Experiment." *The Review of Economic Studies*, 82(2 (291)): 791–824.

Woodford, Michael. 2003. "Imperfect Common Knowledge and the Effects of Monetary Policy." In *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps.* , ed. Philippe Aghion, Roman Frydman, Joseph Stiglitz and Michael Woodford, 25–58. Princeton University Press.

Zafar, Basit. 2011. "Can subjective expectations data be used in choice models? evidence on cognitive biases." *Journal of Applied Econometrics*, 26(3): 520–544.