Overpersistence bias in individual income expectations and its aggregate implications*

Filip Rozsypal[†] Kathrin Schlafmann[‡]

November 11, 2016

Job Market Paper

(latest version available here)

Abstract

We study the role of household income expectations for consumption decisions. Using micro level data, we first document an income-related systematic component in household income forecast errors. We show that these systematic errors can be explained by a modest deviation from rational expectations, where agents overestimate the persistence of their income process. We then study the implications of this overpersistence bias in a quantitative model. Low income households who overestimate the persistence of their income are too pessimistic about their future income. This has two effects. First, they are unwilling to borrow to smooth their consumption even though their borrowing constraint is not binding, thereby allowing the quantitative model to match the distribution of liquid assets across the income distribution. Second, they have lower marginal propensities to consume than their fully rational counterparts. This implies that standard models of household consumption overpredict the effectiveness of government stimulus payments if they do not take deviations from rational income expectations into account.

JEL codes: D14, D84, D91, E21, G02, H31

KEYWORDS: household income expectations, savings, durable consumption

*We would like to thank Vasco Carvalho, Wouter Den Haan, Chryssi Giannitsarou, Per Krusell, Hamish Low and Ricardo Reis as well as seminar participants at London School of Economics, Greater Stockholm Macro Group, University of Cambridge, University of Bonn, CEF Bordeaux, NorMac 2016, CEPR Household Finance 2016 and Workshop on Household Surveys in Macroeconomics at University of Hamburg for many valuable comments and suggestions.

[†]Corresponding Author, Contact Information: Filip Rozsypal, London School of Economics, Centre for Macroeconomics, London, UK, email: f.rozsypal1@lse.ac.uk

[‡]Contact Information: Kathrin Schlafmann, Institute for International Economic Studies (IIES), Stockholm, Sweden, email: kathrin.schlafmann@iies.su.se

1 Introduction

Income risk is 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 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.

This papers makes two contributions. First, we use micro data on household income expectations and provide evidence of non-rationality in the form of systematic biases. We argue that our findings are consistent with households overestimating the persistence of their individual income process and being too pessimistic about the development of the aggregate economy. Second, we show how this bias affects consumption and savings behavior in an otherwise standard model of durable consumption. We find that allowing for the biases identified in micro data helps to improve the fit of the cross-sectional distributions of liquid assets. In particular, this mechanism can explain why low income households do not borrow more to smooth consumption. Moreover, we show that standard models of household consumption overestimate the aggregate effectiveness of government stimulus policies if they do not account for income expectation biases as found in the data.

Using data from the Michigan Survey of Consumers, the first part of the paper shows that current income is systematically correlated with the error people make when they forecast their individual future income growth. In particular, 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. We argue that this pattern is generated by people overestimating the persistence of their income process. We hence call this bias *overpersistence bias*. Moreover, we show that people across the whole income distribution are too pessimistic about aggregate variables such as inflation and the unemployment rate.

In the second part of the paper we use a partial equilibrium model of durable and non-durable consumption choice to analyze how the overpersistence bias and aggregate pessimism affect consumption decisions across the income distribution. We implement the overpersistence bias by allowing the agents' belief about the autocorrelation parameter in their income process to differ from the true underlying parameter. Even though all households share the same 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 realizations 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. These heterogeneous effects allow us to match the observed expectation errors across the whole income distribution with only two parameters: the households' beliefs about the autocorrelation of individual income and a parameter that governs the aggregate pessimism.

We find that biased income expectations have differential effects on the behavior of households depending on their relative position in the income distribution. High income households hold similar assets under biased and under fully rational expectations. The reason is that for them overpersistence bias and aggregate pessimism have opposing effects and cancel each other out. At the same time, the portfolio choice of low income households varies significantly with the accuracy of income expectations. Low income households who have biased expectations are too pessimistic about their future income. This is why they do not want to borrow to smooth consumption even though their borrowing constraint is not binding and they would be able to borrow. We show that this mechanism allows an otherwise standard model to fit the distribution of liquid assets as well as durable holdings across different income groups. The model with fully rational income expectations, on the other hand, would predict counterfactually large amounts of borrowing. Such large fractions of households with negative assets is a common feature of Bewley-type models (see Huggett (1996) and more recently in an overview De Nardi (2015)). Including biases in income expectations as seen in the data allows the model to overcome this counterfactual behavior.

We further investigate how the deviations from rational expectations affect the marginal propensity to consume (MPC). We show that the overpersistence bias reduces the difference between the MPC of low and of high income households relative to the fully rational model to a level in line with empirical estimates (Johnson et al., 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). In both recent recessions of 2001 and 2008 the U.S. government employed this policy by handing out one-off cash transfers. However, assuming a balanced budget and a progressive tax system, such programs redistribute wealth from high to low income agents. Hence the higher is the difference between the MPC of low and high income households, the higher is the aggregate consumption response. The results in the paper reveal that low income households with biased expectations have lower MPCs than their rational expectations counterparts. High income households, on the other hand, turn out to have similar levels of MPC in both expectation scenarios. Standard models which do not take biases in income expectations into account hence overestimate the effectiveness of government stimulus policies.

Finally, we investigate the implications of an alternative way to achieve lower levels of borrowing: tightening the borrowing limit. We find that fully rational agents are particularly responsive to changes in the borrowing constraint. At the same time, however, we document that tight borrowing constraints can strongly inflate the model implied MPC and hence the effectiveness of stimulus policies. Our analyses therefore show that having a realistic mechanism for why people do or do not borrow can have important aggregate consequences.

The paper contributes to the literature in two fields. First, it contributes to the growing literature about expectation formation. Most of this literature has analyzed expectations about aggregate variables, in particular inflation expectations (see, e.g. Carroll (2003) and more recently Malmendier and Nagel (2015)). In contrast, we focus on individual level income expectations and realizations. Household income expectation have hardly been studied in the literature. Dominitz and Manski (1997), Dominitz (1998) and Das and van Soest (1999) are 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. The difference to the current paper is that the Dutch data set asks households only 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. 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. 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 enables us to study the effects of this bias in a fully specified consumption-saving framework.

The second strand of literature that this paper directly contributes to is the literature on durable versus non-durable consumption (see, e.g., Bertola and Caballero (1990), Grossman and Laroque (1990) and Bar-Ilan and Blinder (1992)) and how this relates to marginal propensities to consume. The two most relevant studies for this paper 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. We show that biased and fully rational expectations have different implications for the joint distribution of liquid assets and income and for the effectiveness of stimulus policies.

The paper proceeds as follows. Section 2 empirically documents the systematic bias in household income expectations and argues that the findings are caused by households overestimating the persistence of their income process and aggregate pessimism. Section 3 describes the partial equilibrium model of durable and non-durable consumption and how the empirical findings are incorporated into the theoretical framework. Section 4 details how we quantify the model. Section 5 shows the effects of biased income expectations on the behavior of households in different income groups and how they affect the distribution of MPCs out of transfer payments. Furthermore, the section discusses the implications of borrowing constraints for MPCs. Section 6 concludes.

2 Household Income Expectations

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. After documenting these facts we analyze possible mechanisms which can generate the observed errors in expectations. We will argue that households seem to overestimate the persistence of their income process so that they fail to sufficiently account for mean reversion of their income relative to the cross-section.

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 this subset of households we 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

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

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?

2.1 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 challenge we face is caused by the fact that there is only imperfect overlap between the periods for which households give expectations and for which they report realizations. 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*.

In both the first interview and the re-interview after 6 months households are asked for their income in the previous calendar year. This implies that 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 hence 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. However, 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.[2]

Hence we proceed in two steps. First, we conduct the analysis on July households only as there is the largest overlap. We find significant effects of income in the first interview on the forecasting error (see table 1, column 4). In the second step, we extend our analysis to the rest of the year by exploiting the fact that households interviewed in the second half of a year provide two consecutive years of income realizations. These households thus provide an

[2]In contrast to our study, Souleles (2004) does not consider the implications of the timing of interviews or the imperfect overlap of expectations and realizations.

Figure 1: Timing of Income Realizations versus Expectations

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

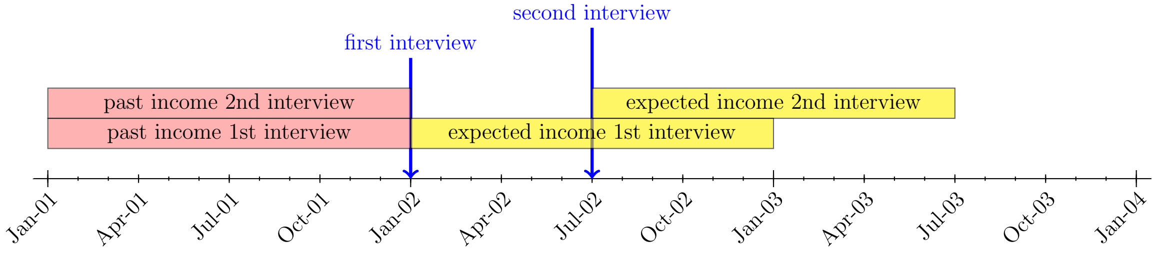

(b) First interview in July 2002 - reported data:

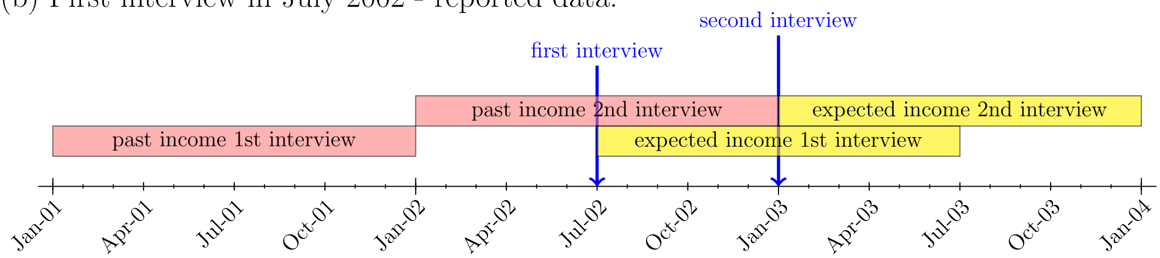

(c) First interview in January 2002 - imputed income:

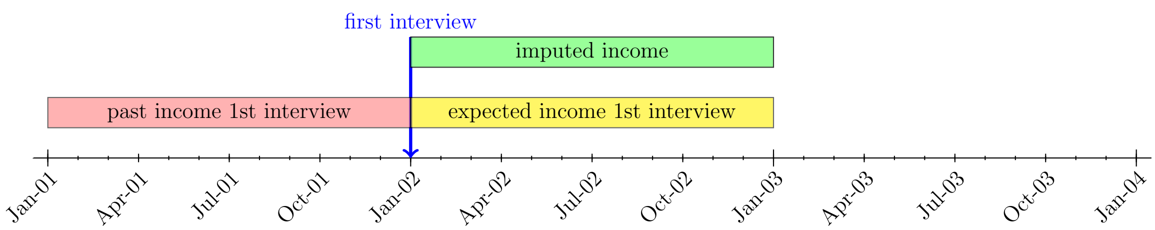

observation of realized income growth that can be used to infer a relationship between income growth in a particular year and the level of income as well as household characteristics in the year prior to that.[3] We use this relationship to impute income growth realizations for the households interviewed in the first half of the year (see panel (c) of figure 1).[4] By imputing we can increase the number of observations as well as improve the timing overlap between expectations and realizations. This is our baseline. Since we find similar results on this full sample as we do on the non-imputed July sample we can be assured that our results are not driven by the imputation procedure. Moreover, we conduct further robustness checks to

[3]We estimate this relationship separately for each year. This specification is therefore fully flexible regarding the effects of aggregate factors in the economy.

[4]A detailed description of the imputation procedure can be found in appendix A.2.

Figure 2: Expectation errors in real income growth

(a) mean (b) mean by income quintile

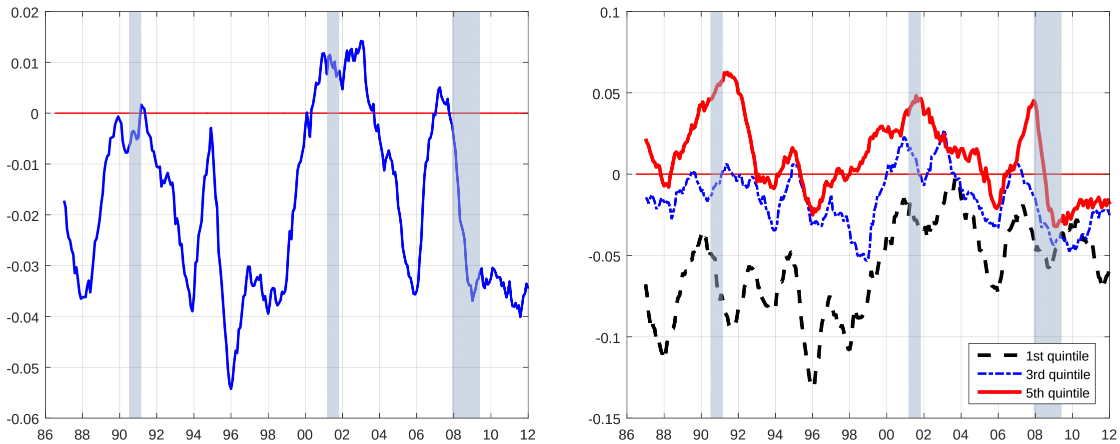

Note: The figure plots the mean expectation errors in individual real income growth smoothed with 12-month moving average filter. Expectation errors are winsorized at 5% and 95%. Data from the Michigan Survey of Consumers and own calculations. Grey areas represent NBER recessions. On the y-axis, 0.01 corresponds to 1 percentage point.

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.

2.2 Expectation Errors Analysis

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.[5] It is negative in most of the analyzed time period (panel (a)). This indicates

[5]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. Appendix A.3 shows the corresponding time series plots to figure 2 for nominal income expectations. Moreover, when we control for household characteristics

that for the population as a whole, people tend to be too pessimistic about their income growth. However, there is considerable heterogeneity in the forecast error by household income. 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).[6] 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.

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.[7] 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), repectively. 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

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.

[6] Households are allocated to income quintiles based on the cross-sectional distribution of per adult income in the year of the first interview.

[7] Appendix A.4 also contains the results when interaction terms of income quintiles with age bins and education dummies are included. Most of these interaction terms are not significant and the relationship between expectation errors and income quintiles is robust to this change: it remains statistically and economically significant and of very similar magnitude as in the main specification.

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

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Note: Table shows regressions results from OLS on 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 included month dummies as additional controls.

Figure 3: Expectation errors in real income by income group

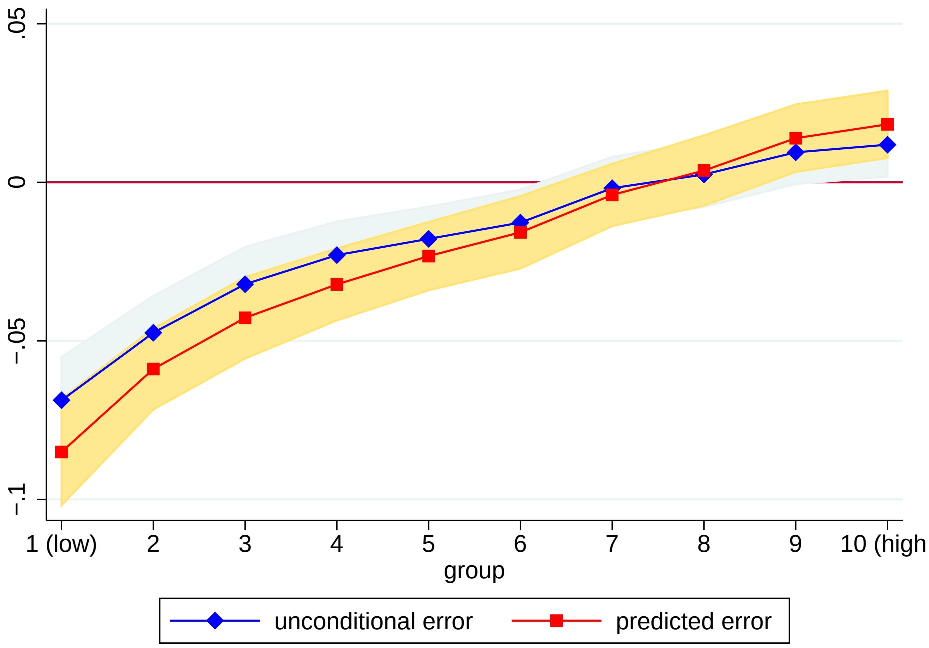

Note: The figure shows the unconditional mean expectation error (blue line, diamonds) and predicted expectation error (red line, squares) in real income growth by income decile. Predicted expectation errors are based on regression results from table 1 column 1, except that income is split in income deciles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

is done on July interviews only using directly reported income changes instead 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.

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. 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

Figure 4: Expectation errors by income group

(a) nominal income growth (b) inflation

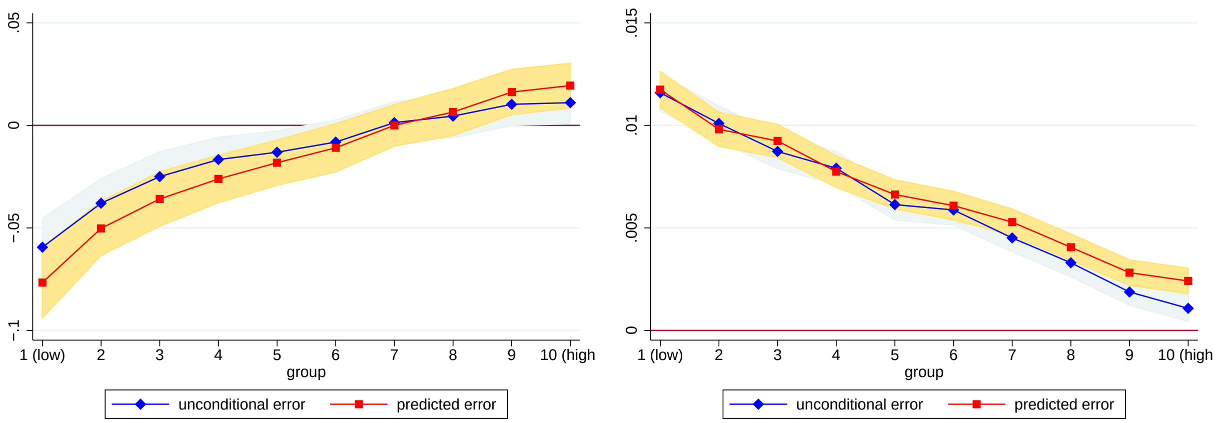

Note: The figure shows the unconditional mean expectation error (blue line, diamonds) and predicted expectation error (red line, squares) by income decile. Predicted expectation errors are based on regression results from table 1 column 5 and 6, except that income is split in income deciles instead of quintiles. Predicted values computed for all other explanatory variables at the weighted sample mean. Bands refer to 95% confidence intervals. On the y-axis, 0.05 corresponds to 5 percentage points.

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.

Another aggregate variable that households in the Michigan Survey of Consumers are asked about is unemployment.[8] 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?*

[8]The survey also elicits expectations about the development of interest rates. Unfortunately, the survey doesn't specify which interest rate, only that people should think of "interest rates for borrowing money". It is hence not clear which interest rates people refer to when they answer the question. This implies that is unclear to which realizations the expectations should be compared.

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

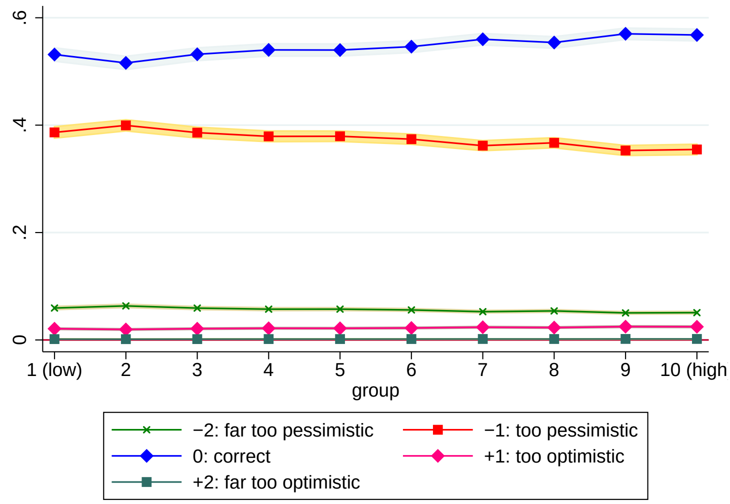

Note: The figure shows the predicted likelihoods of each outcome category of unemployment expectations (-2 (far too pessimistic) to +2 (far too optimistic)) by income decile. Predicted likelihoods are based on a ordered logit regression of categorical forecast errors on income deciles and other demographics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

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.[9] Categorical expectation errors are then defined as "categorical expectation" - "categorical 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).[10] 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.[11]

The analyses in this section thus reveal two forms of bias in household expectations.

[9] 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". We 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.

[10] See appendix A.5 for the full regression results.

[11] This finding of general pessimism in aggregate variables is in line with the results in Bhandari et al. (2016) who show that the average unemployment and inflation expectations are on average too pessimistic across various population groups (including income groups) relative to the Survey of Professional Forecasters.

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 their forecasts of aggregate variables.

2.3 Mechanism: Overestimation of Persistence in Income Process

In this subsection we argue that the observed pattern can be generated by people overestimating the persistence of their income process. In the next subsection we will consider alternative explanations and show that they are not consistent with the empirical findings.

Formally, overestimating the persistence of income can be described as follows. Assume that income (net of age effects and the effects of other demographics) is generated by the process

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

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

where P_{it} is a persistent component and V_{it} is a transitory shock. Persistent income depends on past persistent income and 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) $\exists! \bar{P}$:

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

and vice versa for $P_{it} < \bar{P}$, where $Y_{i,t+1|t}|P_{i,t}$ is the conditional mean of $Y_{i,t+1}$ given $P_{i,t}$ formed under the wrong expectations.

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

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

The theorem thus states that overestimating the persistence of the income process generates expectation errors in income growth that are (a) positive if permanent income is above

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

(a) realized growth (b) growth expectation

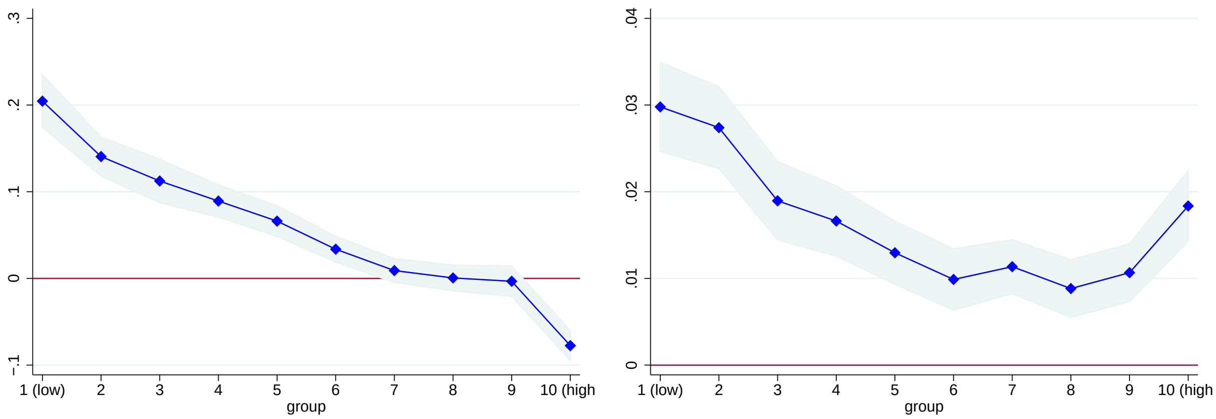

Note: The figure shows the predicted realized growth (panel (a)) and growth expectations (panel (b)) in real income by income decile. Predicted values are based on OLS regression results from regressing individual realized growth rates or expectations on all regressors as in table 1. Sample: for realized 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. On the y-axis, 0.01 corresponds to 1 percentage point.

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. Appendix B contains the proof of the theorem.

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.[12] 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. 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 favour of households overestimating the persistence of their income

[12] These predicted values have been constructed from estimating equation (2) where the outcome variable Z_i is set to the reported realized income growth g_i of households interviewed for the first time in July to December (only those households directly report income changes). Detailed estimation results can be found in appendix A.6.

Figure 7: Expectation errors in real income by age group

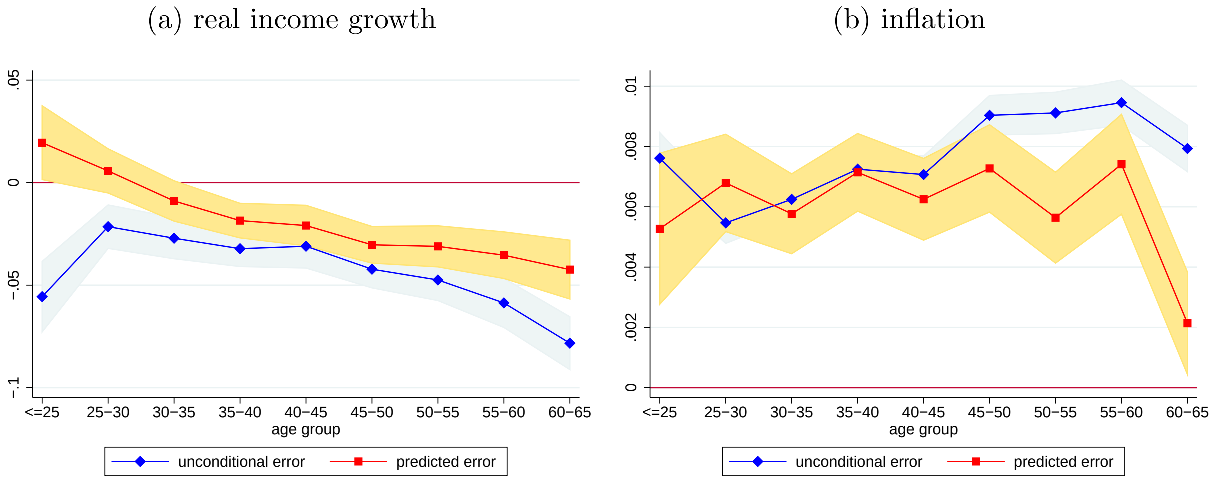

Note: The figure shows the unconditional mean expectation error (blue line, diamonds) and predicted expectation error (red line, squares) by income decile. Predicted expectation errors are based on regression results from table 1 column 1 and 6, except that age is split into 5-year age groups instead of the quadratic term in age. Predicted values computed for all other explanatory variables at the weighted sample mean. Bands refer to 95% confidence intervals. On the y-axis, 0.01 corresponds to 1 percentage point.

process.

2.4 Alternative Mechanisms?

In this subsection we go through alternative mechanisms that could potentially generate the same pattern of expectation errors. We will argue that none of them are consistent with the empirical results.

Learning One potential explanation could be that people need to learn about their income potential over time, so that young households could be expected to make larger errors than older households. While in the above regressions we already control for age effects, it might still be the case that expectation errors vary systematically with age. Figure 7 shows the unconditional as well as the predicted expectation errors for different age groups (holding all other characteristics, including income, at their sample mean). Panel (a) shows that the unconditional mean error is hump-shaped in age. However, once all other characteristics are controlled for, expectation errors are in fact decreasing with age, indicating that people become more and more pessimistic with age. It is not the case that expectations would improve as households age. Moreover, panel (b) shows that there is no clear pattern in inflation expectations with regards to age. Based on this result we conclude that people do not seem to learn about their income potential over time.

Table 2: Effect of Recent Experience on Growth Expectations

	(1) real	(2) real	(3) nominal	(4) nominal
past expectation	0.372***	0.374***	0.373***	0.374***
	(0.016)	(0.016)	(0.016)	(0.016)
past realized growth		−0.021***		−0.022***
		(0.004)		(0.004)
Income Quintile				
1st	0.004	0.007	0.007	0.009**
	(0.004)	(0.004)	(0.004)	(0.004)
2nd	0.002	0.003	0.004	0.005
	(0.004)	(0.004)	(0.004)	(0.004)
4th	−0.005	−0.006*	−0.005	−0.006*
	(0.004)	(0.004)	(0.003)	(0.003)
5th	−0.008**	−0.010**	−0.008**	−0.010**
	(0.004)	(0.004)	(0.004)	(0.004)
Constant	0.061***	0.059***	0.070***	0.068***
	(0.022)	(0.022)	(0.022)	(0.021)
Observations	15931	15931	17210	17210
R^2	0.185	0.187	0.182	0.184

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Note: OLS estimation of individual growth expectations in 2nd interview as a function of past expectations and recent experience; estimation on sample 2HP (households with first interview in 2nd half of year and reinterview). Additional (unreported) control variables the same as in previous regressions: education, age, age^2, racial background, number of adults, gender, marriage status, region and time dummies.

Inability to distinguish between permanent and transitory shocks In the income process there are two types of idiosyncratic shocks which differ in their persistence. The first type of shock is persistent. The other type is completely transitory. Could an inability to distinguish between the two shocks generate the pattern of expectation errors that we observe in the data? If households cannot tell the shocks apart and observe only overall income, they have to rely on some form of filtering to form beliefs about the current state. From linear projection theory we know that Kalman filtering is unbiased and optimal for linear systems and normal shocks. Hence there cannot be a systematic error conditional on past income developments if people form their beliefs optimally. A sketch of a formal proof can be found in appendix C.

Extrapolation of recent experience One explanation why current income can predict expectations about future income growth could be that people overweigh their recent experience. This would imply that households with a recent increase in income - which is correlated with being in a higher income group, all else equal - would expect another increase in the

future. We test for this explanation by regressing the growth expectations in the second interview on past expectations and recent experience (as well as on the other control variables we included in previous regressions). Table 2 shows that past expectations explain a large portion of current expectations, which means there is persistence in expectations on the individual level. Recent experience, on the other hand, turns out to be significantly negative. This shows that households do not extrapolate from their recent experience. In fact, they seem to anticipate that there is mean reversion in their income process. Note, however, that the magnitude of this anticipated reversion is economically small. We can hence exclude extrapolation from recent experience as an explanation of the systematic expectation errors by income groups.

Systematically wrong expectations about aggregates Another explanation for the observed pattern in expectation errors could be that households have biased expectations about aggregate conditions that vary systematically with their relative position in the income distribution. However, as seen in the analyses above, household expectations about aggregate variables - such as inflation and the unemployment rate - are too pessimistic across the whole income distribution. Moreover, the magnitude of this bias doesn't vary much with income groups. Expectation errors in aggregate variables thus cannot explain the shift from overpessimism to overoptimism we observe as we move along the income distribution.

Other mechanism Brunnermeier and Parker (2005) describe a setting where agents find it optimal to have too optimistic expectations. Alternatively, it might be possible that in order to attempt high risk-high reward projects, one needs to underestimate the chances of failure. The overoptimism for high income households could then arise as a result of survival bias. However, neither of these mechanism can explain why low income households are on average too pessimistic in their expectations. Regarding the low income agents, it might be possible that it is very costly to face unusually bad realizations. This could explain why they might find it optimal to behave as if they expected the worst possible outcome. Such agents might also optimally form overly pessimistic expectations to avoid disappointment. However, this mechanism cannot explain the overoptimism for high income households.

To sum up, from the analyses in this section 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, but not with alternative explanations. Second, households across the whole income distribution are too pessimistic about aggregate conditions. In the remainder of the paper we will thus

analyze how overestimation of persistence in income in combination with aggregate pessimism affects household consumption choices. Moreover, we will show the implication of these biases in expectations on the the marginal propensity to consume across the income distribution and ultimately on the effectiveness of fiscal transfers.

3 Model of Household Consumption Choices

In this section we analyze how the biases in income expectation that we documented in the empirical part affect consumption and saving decisions. We build a quantitative framework where we compare the behavior under rational expectations with the behavior under biased expectations. The model setting is close to the one used by Berger and Vavra (2015). Apart from the possibility of biased income expectation the most important difference is in the treatment of the borrowing constraint. Whereas Berger and Vavra assume that agents can only save (no borrowing), we allow households to borrow up to a limit determined by its income state and durable holdings.[13] This assumption is not only more realistic, but it also has important consequences. First, a significant fraction of US households holds negative liquid assets. In order for the model to fit the data borrowing is hence essential. At the same time, however, Bewley-type models typically generate too many households with zero or negative assets compared to the data, in particular if borrowing is allowed (see, e.g., Huggett (1996) and De Nardi (2015)). We will show that including biased income expectations as seen in the data can overcome this counterfactual prediction of standard models. Lastly, the ability to borrow, other things equal, reduces the number of constrained agents and consequently affects the marginal propensity to consume.

3.1 Household Optimization Problem

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.[14] The stock of durable goods depreciates and is subject to 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

[13]Kaplan and Violante (2014) allow for borrowing, but their borrowing limit is independent of the value of the durable good. The main difference between our setting and Kaplan and Violante (2014) is that the latter analyzes a life-cycle model, whereas we have an infinite horizon setup (which we share with Berger and Vavra (2015)).

[14]Appendix D shows the results of a version of the model without durable goods. The results of the full model hold in this restricted setting. As is to be expected, however, this simplified model is not able to accurately capture the cross-sectional distribution of assets.

income.

Households maximize their discounted life time utility[15]

$$\max_{\{c_t\}_{t=0}^{\infty},\{d_t\}_{t=0}^{\infty},\{s_t\}_{t=0}^{\infty}} \quad \sum_{t=0}^{\infty} \beta^t \ \mathsf{E}\big[U(c_t, d_t)\big], \tag{6}$$

subject to the following budget constraint

$$c_t + d_t + s_t + A(d_t, d_{t-1}) \leq R(s_{t-1}) + Y_t + (1-\delta)d_{t-1}. \tag{7}$$

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

$$R(s_t) = [1 + r\,(s_t)]s_t \text{ where } r(s_t) = \begin{cases} r^l & \text{if } \ s_t > 0 \\ r^b & \text{if } \ -(\kappa_y P_t + \kappa_v d_t) \leq s_t \leq 0 \end{cases} \tag{8}$$

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 $\kappa_v d_t$).

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 $A(d_t, d_{t-1})$:

$$A(d_t, d_{t-1}) = \begin{cases} 0 & \text{if } d_t = (1-\delta)d_{t-1} \\ F^d(1-\delta)d_{t-1} & \text{otherwise.} \end{cases} \tag{9}$$

Equation (9) 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

$$U(c, d) = \frac{\left[\left((1-\theta)c^{\frac{\xi-1}{\xi}} + \theta(\bar{d} + d)^{\frac{\xi-1}{\xi}}\right)^{\frac{\xi}{\xi-1}}\right]^{1-\gamma}}{1-\gamma}. \tag{10}$$

Note that every household obtains utility from a small free stock of durable $\bar{d}$. This captures the fact that it is not necessary for a household to own a car, furniture or a washing machine.

[15]To simplify notation we have dropped the subscript i which indicates the individual household. We will only introduce it when describing the income process where it is important to differentiate between aggregate and idiosyncratic shocks.

This specification of the utility function hence enables the model to allow for small levels of durable stocks.

3.2 Income Process and Income Expectations

The only source of risk in the model is income risk. Household i's income at time t is the result of 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 $T_{i,t}$:

$$Y_{i,t} = Z_t \cdot P_{i,t} \cdot T_{i,t}. \tag{11}$$

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

$$T_{i,t} \sim \log N\left(-\sigma_T^2/2, \sigma_T^2\right). \tag{12}$$

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

$$\log P_{i,t} = \rho \log P_{i,t-1} + \epsilon_{i,t}^P, \qquad \epsilon_{i,t}^P \sim N(0, \sigma_P^2) \tag{13}$$

and the aggregate persistent component is a two state Markov process

$$\mathbb{Z} = \begin{bmatrix} Z^h \\ Z^l \end{bmatrix}, \qquad \Pi_Z = \begin{bmatrix} \pi_{11} & 1-\pi_{11} \\ 1-\pi_{22} & \pi_{22} \end{bmatrix}, \tag{14}$$

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

Motivated by our findings discussed in the previous section, we now introduce the two expectation biases. The overpersistence bias in expectations is implemented by letting agents believe that the persistence of the idiosyncratic component P is higher than its true value. Formally, agents believe that their persistent income component evolves according to the following process:

$$\log P_{i,t} = \hat{\rho} \log P_{i,t-1} + \epsilon_{i,t}^P, \qquad \epsilon_{i,t}^P \sim N(0, \sigma_P^2), \tag{15}$$

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

The overpessimism 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:

$$\hat{Z}_{t+1} = \mu \mathsf{E} Z_{t+1} = \mu \Pi_Z Z_t. \tag{16}$$

In the calibration both bias parameters - the overpersistence belief $\hat{\rho}$ and the overpessimism parameter μ - will be found by matching the empirically observed forecasting errors by

income quintile with the ones generated in the model.

4 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 (income process, 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 parameters (preferences and beliefs) to match moments in the data. Within this calibration step we first estimate the belief parameters which govern the bias in household income expectations to match the expectation errors by income quintile that we observe in the data.[16] Once we have obtained these belief parameters we calibrate the remaining preference parameters to fit the empirical distributions of liquid assets and durable holdings. Table 3 reports the resulting parametrization.

4.1 Exogenous Parameters of the Environment

Technology 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 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 et al. (2008) for the Consumer Expenditure Survey.

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

[16] This part only depends on the specification of the model income process but is independent of all other parameters.

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	0.6
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

can borrow up to 60% against the value of their durable and set $\kappa_v = 0.6$.

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 has been owned by 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 $\delta = 0.05$.

Income process For the parametrization of the income process we follow Storesletten et al. (2004) who estimate an income process with persistent and idiosyncratic shocks. We transform their income process to quarterly frequency and obtain the following parameters:

Table 4: Mean expectation errors

	data	model
income quintile 1	-0.072	-0.068
income quintile 2	-0.037	-0.040
income quintile 3	-0.019	-0.021
income quintile 4	-0.000	-0.004
income quintile 5	0.016	0.020

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.

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). 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.[17]

4.2 Parameters Estimated to Match Data Moments

There are seven remaining parameters: the two parameters that govern the magnitude of the bias in income expectations (overpersistence bias $\hat{\rho}$ and aggregate pessimism μ) and five preference parameters that shape the utility function and time preference. We are able to estimate these two sets of parameters sequentially since the belief parameters are chosen to match observed errors in income expectations which are independent of preference parameters. We therefore first determine the magnitude of the bias in income expectations and then take these parameter values as given when we search for the remaining preference

[17] The exact formula is

$$\text{avg\_dev} = \frac{1}{T_{pos}} \sum_{t=1}^{T} \hat{y}_t \cdot I(\hat{y}_t > 0) - \frac{1}{T_{neg}} \sum_{t=1}^{T} \hat{y}_t \cdot I(\hat{y}_t < 0) \tag{17}$$

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 5: Model fit

		quantile							mode
		0.05	0.10	0.25	0.50	0.75	0.80	0.90	
liquid assets	data	-1.29	-0.88	-0.30	0.03	0.76	1.36	5.46	-0.02
	model	-1.04	-0.79	-0.37	-0.05	0.10	0.14	0.25	0.01
durables	data	0.13	0.20	0.39	0.79	1.43	1.62	2.21	0.23
	model	0.24	0.30	0.44	0.72	1.16	1.29	1.65	0.40

Note: Selected moments generated by the model compared to Survey of Consumer Finances.

parameters.

Beliefs about the income process 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 4 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. The choice of these parameters depends only on the parameters of the income process and is hence fully independent of all the other parameters in the model.

Preference parameters The remaining five parameters are preference parameters which affect the allocation 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. 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 data distributions we target have been obtained from the Survey of Consumer Finances (SCF), waves 1992-2010. Liquid assets are defined as checking accounts, savings accounts, stocks, bonds, and mutual funds minus outstanding credit card debt after last payment and outstanding auto loans. The data counterpart for durable goods is 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.[18]

Formally, we choose the parameters which minimize the difference between the quantile

[18] Households without any vehicle constitute 13% of the sample population.

Figure 8: Model fit

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

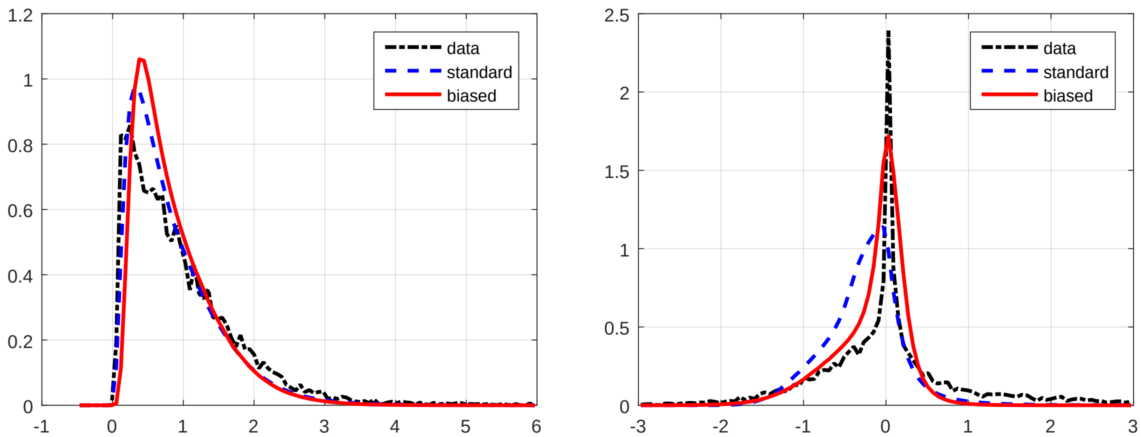

Note: The figure depicts the distribution for (a) durable goods and (b) liquid savings. Data distributions (dash-dotted black line) are compared to the distributions implied by model which allows for biased expectation (solid red line) and the model where expectations are assumed to be rational (dashed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

function of the model with biased beliefs and the data:

$$\Omega = \arg\min \sum_{X} \left(\sum_{i}^{\bar{i}_X} |q_i(X) - \tilde{q}_i(X)| \right) \tag{18}$$

where $\Omega = \{\beta, \gamma, \theta, \xi, \bar{d}\}$, q_i and $\tilde{q}_i$ is i-quantile of the data and the model and X represents either liquid savings or durable stocks. We do not include the top 20% of the distribution for liquid assets and the top 5% of durable stocks in our objective function ($\bar{i}_s = 80$ and $\bar{i}_d = 95$). Under any parametrization, the model struggles to replicate the thick right tails, in particular in liquid assets. Our definition of liquid assets includes stocks and bonds, which people often use for longer term saving since they offer higher but uncertain returns. The model does not capture these life cycle motives for saving. Agents hold liquid assets for transactionary (due to the adjustment costs in durables) and precautionary reasons. Nevertheless, since stocks and bonds are highly liquid, we include them in the definition of liquid assets.

The resulting parameter 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 $\bar{d} = 0.5$. Table 5 and figure 8 show the resulting model fit.

5 Effects of Income Expectations on Household Consumption Behavior

In this section we first show how the beliefs about income expectations affect the behavior of households in different income groups and show that it is in line with the empirical distributions. We demonstrate that under rational expectations the model predicts counterfactually large borrowing for low income households. Allowing for overpersistence belief and aggregate pessimism in income expectations hence reconciles the model predictions with the data. Furthermore, we show how biased income expectations affect the marginal propensity to consume (MPC) out of unanticipated transfer payments. We find that the overpersistence bias differentially affects the MPC in different income quintiles. Low income households turn out to have a smaller MPC if they have biased expectations while the MPC of high income households is hardly affected by the beliefs. Overall, the differences in MPC's across the income distribution are hence smaller than what would be predicted under rational expectations.[19]

5.1 Effects on Behavior Across Different Income Groups

Figure 9 shows the distribution of durable goods for households in the lowest and highest income quintiles. The black dash-dotted line refers to the data distribution while the solid red line depicts the distribution in the model where households have biased income expectations. The model is able to match the cross-sectional variation in durable holdings well. Moreover, the figure also plots the distribution implied by the model if households had rational expectations (blue dashed line). In terms of durable holdings, biased expectations 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 10 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 low income group depends on what households believe about their future income. Low income households with biased beliefs are too pessimistic and expect their income to decrease in the future. They are therefore less willing to borrow even though their borrowing constraint is not binding. Figure 10(a) shows that this mechanism allows the model with biased income expectations to fit the empirical distribution of liquid assets in the lowest

[19]In this section we compare the implications of a model with biased expectations to the implications of the same model (i.e. same parametrization) under rational expectations. In appendix E we show the corresponding results when we instead calibrate the parameters to maximize the fit of the fully rational model. The qualitative results are the same as what is described in the main text.

Figure 9: Durable stock by income quintile (d)

(a) lowest quintile (b) highest quintile

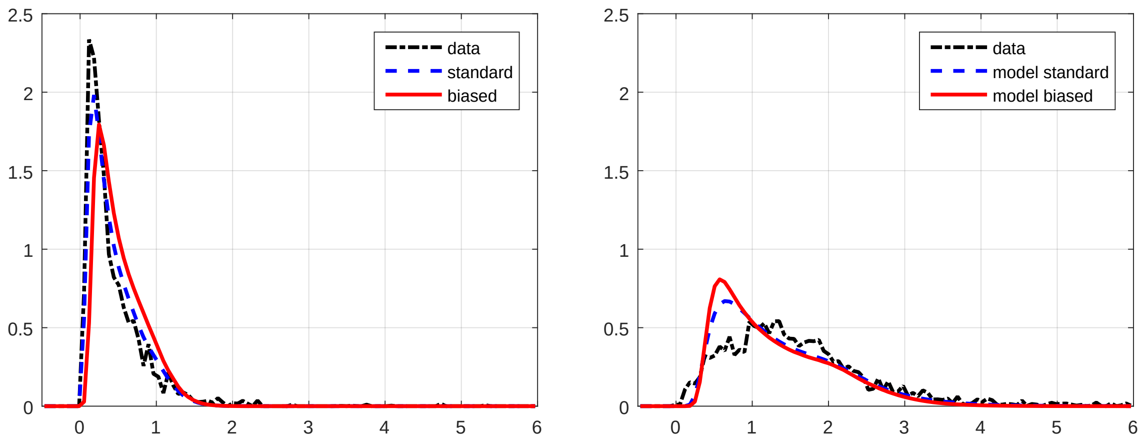

Note: The figure depicts the distribution of durable goods in the model versus data for different income quintiles. The panels show the data distribution (dash-dotted black line) against the model distribution when households have non-rational expectations (solid red line). For comparison, the distribution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

Figure 10: Liquid assets by income quintile (s)

(a) 1st, rational versus both bias (b) 5th, rational versus both bias

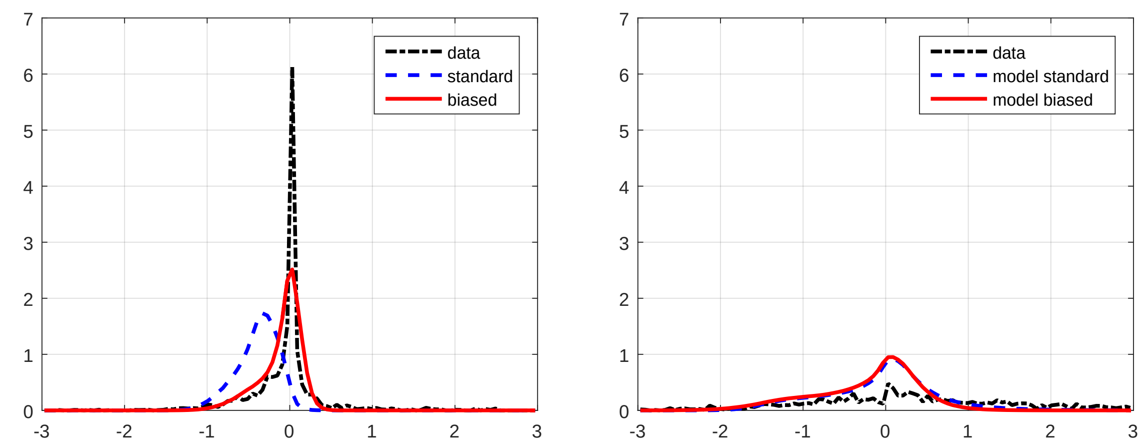

Note: The figure depicts the distribution of liquid assets in the model versus data for different income quintiles. The panels show the data distribution (dash-dotted black line) against the model distribution when households have non-rational expectations (solid red line). For comparison, the distribution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

income group very well. If people had rational expectation instead, the model would predict counterfactually large amounts of borrowing (mode of -0.5 versus 0 in the data).

5.2 Implications for Marginal Propensity to Consume

Government stimulus 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. However, how effective these stimulus programs are depends on how much households effectively spend out of the transfer. Moreover, 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 ratio between the MPC of low income households relative to the MPC of high income households. This ratio can be seen as a measure for the first order effect in the transfer multiplier. In this section we show that biased income expectations directly affect this measure and help realign the model implied ratio with the one found in empirical studies of the last two stimulus payments in the U.S.

Figure 11 shows the reaction of households to a one-time, unanticipated transfer payment of varying size. Panel (a) shows the fraction of the transfer payment that the population as a whole spends on non-durable consumption. The solid red line shows the behavior of households with biased beliefs while the blue dashed line depicts the MPC if households instead had rational beliefs. The dash-dotted magenta line further shows the MPC of households with biased beliefs if they had the same level of savings and durable stock as their rational agents counterpart. Panels (b) and (c) show the corresponding MPCs of the lowest and highest income quintile, respectively. Low income households with biased expectations have an average MPC that is between 3-4 percentage points lower than the MPC of rational households, depending on the size of the transfer. High income households, on the other hand, spend about the same fraction of the transfer payment whether they have biased expectations or not. Moreover, the figure shows that the difference in MPCs between households with biased expectations and rational expectations is mostly driven by the difference in the level of assets households hold. If households with biased beliefs had the same portfolio as rational households, the differences in MPCs between the two types would be much smaller. This smaller reduction in MPCs is the direct effect of biased expectations. The remaining larger part of the difference is driven by the different asset positions that households with biased expectations accumulate over time. They are less likely than their rational expectations counterparts to get close to the borrowing constraint and hence have a lower MPC.

Overall, there are two consequences of biased income expectations on MPCs. First, the

Figure 11: MPC out of unexpected transfer

(a) aggregate (b) lowest income (c) highest income

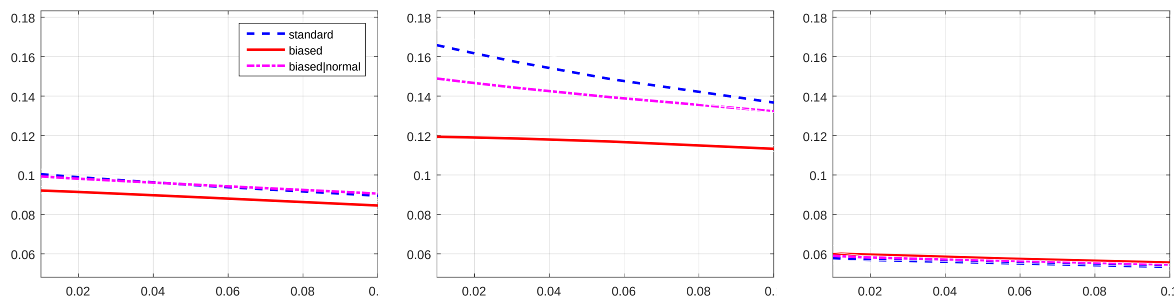

Note: The figure depicts the fraction of an unanticipated one-time transfer payment of varying sizes that is spent on non-durable consumption under different expectation scenarios: the red line depicts the MPC under biased expectations, the dashed blue dashed line depicts the MPC under rational expectations 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 standard agents. Panel (a) shows the MPC in the aggregate population while panels (b) and (c) show the MPC for the lowest and highest income quintile. Transfer sizes are expressed as fractions of average quarterly income in the economy. MPC is computed as average increase in consumption relative to the size of the transfer.

MPC in the population as a whole is lower under biased expectations than what would be predicted if people had rational expectations. Second, the differences in marginal propensities to consume between different income groups are smaller if we take non-standard expectations into account.

Table 6 further compares the model implied MPCs under both expectation scenarios with the MPCs that were found in the literature for the two U.S. stimulus payments in 2001 and 2006.[20] The first observation is that the level of MPCs in the model is lower than the estimates obtained for the two recent stimulus payments. However, the second observation concerns the ratio between the MPC of low income households relative to the one of high income households, i.e. our measure of the first order effect in the transfer multiplier. According to Johnson et al. (2006) and Parker et al. (2013), in 2001 that ratio was equal to 2.33 which was much larger than in 2008 where the MPC of the two groups were much more similar with a ratio of 1.16. The ratio implied by the model that allows for biased expectations generates a ratio of 1.88 which is well between the two estimates. Under the assumption of rational expectations, however, the model predicts a ratio of low income MPC to high income MPC that exceeds even the larger empirical estimate from the 2001 tax rebate. The model with rational expectations therefore overpredicts the effectiveness of the stimulus

[20] The data estimates for 2001 are taken from Johnson et al. (2006) who analyze the 2001 tax rebate of \$300-\$600 per adult (income groups defined as: low < \$34K, high > \$69K). The data estimates for 2008 are taken from Parker et al. (2013) who analyze the economic stimulus in 2008 of \$300-\$600 per adult and \$300 per child (income groups defined as: low < \$32K, high > \$75K).

Table 6: MPC: model versus data

	model		data	
	biased beliefs	rational beliefs	stimulus 2001	stimulus 2008
MPC low income	0.12	0.16	0.60	0.24
MPC high income	0.06	0.06	0.26	0.21
low/high	1.88	2.61	2.33	1.16

Note: The table compares the marginal propensity to consume out of an unanticipated transfer payment in the model to estimates based on the stimulus payments in the U.S. in 2001 and 2008. The estimates for the year 2001 are taken from Johnson et al. (2006) while the estimates for the year 2008 are taken from Parker et al. (2013).

payment. Allowing for biased income expectations as documented in the data, however, reconciles the model implied cross-sectional variation in MPCs with empirical estimates. It is therefore important to take income expectations into account when considering stimulus policies as a means to increase aggregate consumption.

5.3 Interaction with Borrowing Constraints

It is well known that Bewley-type models tend to imply too many households with zero or negative assets (De Nardi, 2015). The model presented in this paper does not have this problem as the households with low income choose not to borrow due to their pessimistic income expectations. One mechanism used in the literature to prevent people from borrowing too much is to impose exogenous borrowing constraints which turn out to be binding. In this section we argue that the setup of the borrowing constraint can be consequential not only for the distribution of liquid assets, but also for the consumption behavior of households.

We discuss in detail two alternative specifications for the borrowing limit: In the first economy households cannot borrow at all (zero borrowing economy). In the second economy households are allowed to borrow up to the median annual income in the population independent of their own current income (generous limit economy). Formally, we replace the borrowing limit $\kappa_y P_t + \kappa_v d_t$ in equation (8) with a constant $\underline{s}$ and solve the model for values of $\underline{s}$ between 0 (zero borrowing) and -4 (four times the quarterly median income). We discuss the results for $\underline{s} \in \{0, 4\}$ in detail and then show the results for the average MPC as a function of $\underline{s}$.

Effects on the distribution of liquid asset Figure 12 shows the distributions of liquid assets in the two extreme economies for households with rational and biased expectations. Naturally, imposing a zero borrowing constraint renders the model unable to fit the fraction

Figure 12: Distribution of liquid assets under different borrowing constraints

(a) no borrowing (b) limit = median annual income

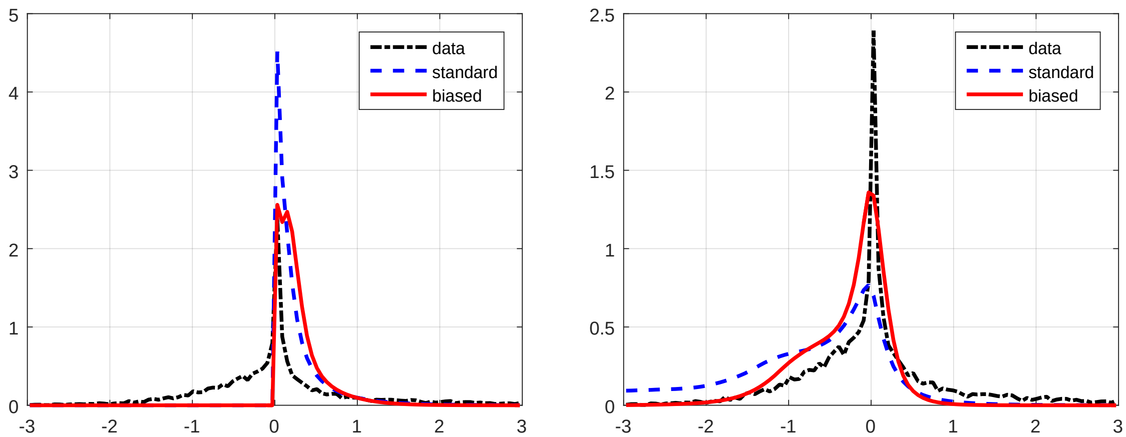

Note: The figure depicts the distribution for liquid assets under two alternative specification for the borrowing constraint: (a) no borrowing at all and (b) unconditional borrowing limit equal to median annual income. Data distributions (dashed black line) are compared to the distributions implied by the 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.

of households that hold negative assets. Instead, a large fraction of households turns out to be constrained at zero liquid assets so that this constraint is binding. This is particularly true for households with rational expectations. With biased expectations, on the other hand, more households decide to build up a small stock of savings. Overall, however, the difference between the two expectation scenarios are smaller in the zero borrowing economy than in the benchmark economy.

In contrast, the behavior of rational agents and of those with biased beliefs differs much more in the generous limit economy. Households with rational beliefs will often make use of the opportunity to borrow so that a significant share of people borrows, even for larger amounts of borrowing. Households with biased beliefs, however, borrow much less. In particular, figure 12(b) shows that they do not even get close to the borrowing limit (which is equal to -4, i.e. four times the median quarterly income). To summarize, relaxing the credit limit leads both rational and biased households to borrow more. However, this effect is much more pronounced for households with rational income expectations. Households with biased income expectations do not want to borrow even though their budget constraint is not binding.

Implications for consumption Changing the borrowing limit also affects the MPC of households. Figure 13 shows the average MPC out of an unexpected transfer payment for

Figure 13: MPC out of unexpected transfer under different borrowing constraints

(a) no borrowing (b) limit = average annual income

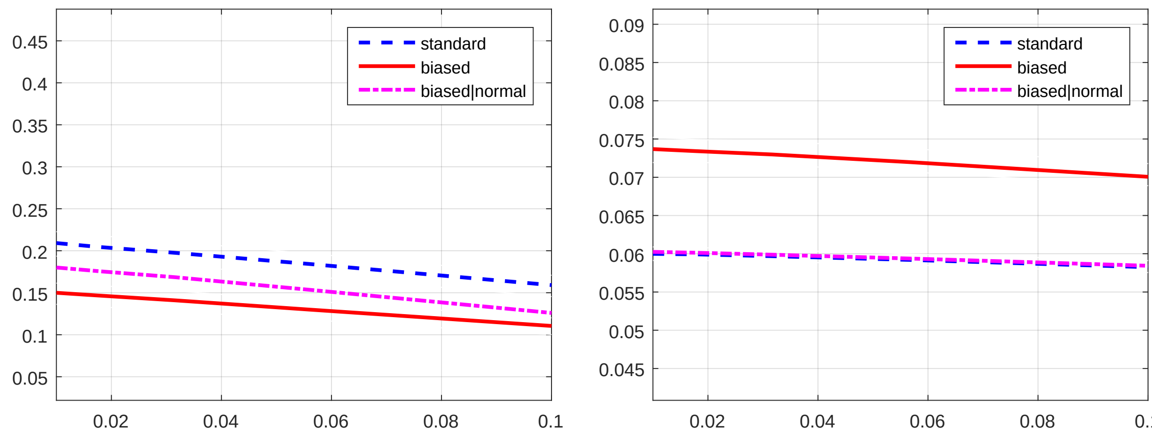

Note: The figure depicts the fraction of an unanticipated one-time transfer payment of varying sizes that is spent on non-durable consumption under different expectation scenarios. X-axis shows the size of the transfer as a fraction of quarterly median income. The red line depicts the MPC under non-rational expectations (overpersistence belief and aggregate pessimism), the blue 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 standard agents. Shown is the average MPC in the population. Transfer sizes are expressed as fractions of average quarterly income in the economy.

the two alternative economies. In the zero borrowing economy the average MPC is almost twice as large as in the benchmark economy. Moreover, the absolute difference between the average MPC of the two types of agents is also larger: households with biased income expectations have an average MPC that is around 5 percentage points lower than that of rational households.

In the generous limit economy the average MPC of both types of agents is lower compared to the benchmark model. This effect is stronger, however, for rational agents. The distributions of liquid assets showed that rational agents make much more use of the availability of credit than biased agents. Relaxing their borrowing constraints therefore reduces the number of constrained agents. Since the borrowing limit in this economy is very generous, the MPC of rational households drops by almost one half. In contrast, households with biased expectations do not want to borrow as much even if they are allowed to. Relaxing their constraints thus has a smaller effect on their behavior and MPC. Figure 13(b) shows that for this strong relaxation of borrowing constraints the MPC of rational agents in fact falls below that of biased agents.

In addition to the two extreme settings for the borrowing constraint $\underline{s}$ that we have discussed so far, we have also solved the model with intermediate levels for the borrowing con-

Figure 14: MPC out of unexpected transfer as a function of borrowing limit

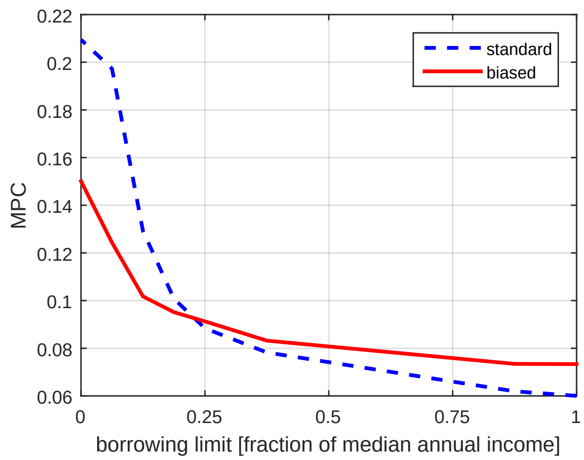

Note: The figure depicts the relationship between the MPC and the borrowing limit. The size of the transfer is 1% of median income and the borrowing limit is on a grid from zero to a full median annual income. The results depicted in figure 13 thus correspond to 0 (no borrowing) and 1.

straint. The resulting average MPCs for the two types of agents are plotted in figure 14. The results suggest that the model predicted MPC can be very sensitive to the borrowing limit, especially for small values of $\underline{s}$. Tightening the borrowing constraint increases the MPC, in particular when the model does not allow for biased income expectations. Admittedly, the quantitative results in this section are model and parametrization specific. Nevertheless, we interpret this observation as a sign to be cautious when setting the parameters regarding borrowing constraints. Tightening the borrowing constraint helps to avoid large amounts of borrowing in a model with fully rational agents. At the same time, however, it can have strong effects on the consumption behavior that the model predicts.

6 Conclusion

In this paper we investigate the role of income expectations on consumption behavior of households. We document a systematic bias in income expectation formation and show its implications for consumption-saving decisions in a quantitative model.

Using household level data from the Michigan Survey 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 over-persistence 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 (2012, chapter 17)). However, to the best 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 find that income expectation biases of the magnitude seen in the data significantly affect the distribution of liquid assets in the cross section. While 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 is not true for low income households. Low income households with biased beliefs are too pessimistic about their future income and are hence unwilling to borrow to smooth consumption. This prediction of the model with biased beliefs is in line with the distribution of liquid assets in the data. If we instead assumed households to have rational expectations the model would predict counterfactually large amounts of borrowing for low income households and for the population as a whole.

The paper further shows that accounting for income expectations is crucial when analyzing the effectiveness of stimulus payments. In the model with rational expectations, the MPC of low income households is too high relative to the MPC of high income households to be consistent with empirical estimates. On the other hand, allowing for biases in income expectations of the magnitude seen in the data leads to a model prediction of this ratio that is well within the range of values estimated for the stimulus payments in the U.S. in 2001 and 2008. If stimulus payments are financed through taxes (which are predominantly paid by high income households), stimulus payments are a form of redistribution. In this light the ratio between the MPC of low income households and high income households can be regarded as a measure of the first order effect in the transfer multiplier. Based on the present analysis we hence conclude that taking biases in income expectations into account is crucial when considering the use of stimulus payments.

We believe that our empirical finding opens an avenue for further research in two main areas. First, while the available data from the Michigan Survey 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 Survey of Consumers we are therefore unable to investigate in detail the process of expectation formation and expectation updating. To learn more about how income expectations are formed it thus seems very important to collect data on income expectations and corresponding realizations in a panel survey.

Second, our analysis suggests 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

Attanasio, O. P., Goldberg, P. K., and Kyriazidou, E. (2008). Credit constraints in the market for consumer durables: Evidence from micro data on car loans*. *International Economic Review*, 49(2):401–436.

Bar-Ilan, A. and Blinder, A. S. (1992). Consumer durables: Evidence on the optimality of usually doing nothing. *Journal of Money, Credit and Banking*, 24(2):258–272.

Berger, D. and Vavra, J. (2015). Consumption dynamics during recessions. *Econometrica*, 83(1):101–154.

Bertola, G. and Caballero, R. J. (1990). Kinked adjustment costs and aggregate dynamics. *NBER Macroecomics Annual*, 5:237–296.

Bhandari, A., Borovicka, J., and Ho, P. (2016). Identifying ambiguity shocks in business cycle models using survey data.

Brunnermeier, M. K. and Parker, J. A. (2005). Optimal expectations. *American Economic Review*, 95(4):1092–1118.

Carroll, C. D. (2003). Macroeconomic expectations of households and professional forecasters. *The Quarterly Journal of Economics*, 118(1):269–298.

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

De Nardi, M. (2015). Quantitative models of wealth inequality: A survey. *NBER Working Paper 21106*.

Dominitz, J. (1998). Earnings expectations, revisions, and realizations. *The Review of Economics and Statistics*, 80(3):374–388.

Dominitz, J. and Manski, C. (1997). Using expectations data to study subjective income expectations. *Journal of the American Statistical Association*, 92:855–867.

Grossman, S. J. and Laroque, G. (1990). Asset pricing and optimal portfolio choice in the presence of illiquid durable consumption goods. *Econometrica*, 58(1):25–51.

Huggett, M. (1996). Wealth distribution in life-cycle economies. *Journal of Monetary Economics*, 38(3):469 – 494.

Johnson, D. S., Parker, J. A., and Souleles, N. S. (2006). Household expenditure and the income tax rebates of 2001. *American Economic Review*, 96(5):1589–1610.

Kahneman, D. (2012). *Thinking, Fast and Slow*. Penguin.

Kaplan, G. and Violante, G. L. (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4):1199–1239.

Malmendier, U. and Nagel, S. (2015). Learning from inflation experiences*. *The Quarterly Journal of Economics*.

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

Parker, J. A., Souleles, N. S., Johnson, D. S., and McClelland, R. (2013). Consumer Spending and the Economic Stimulus Payments of 2008. *American Economic Review*, 103(6):2530–53.

Souleles, N. 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, K., Telmer, C. I., and Yaron, A. (2004). Cyclical Dynamics in Idiosyncratic Labor Market Risk. *Journal of Political Economy*, 112(3):695–717.

A Further Details about Empirical Analyses

A.1 Sample Selection

The Michigan Survey of Consumers interviews around 500 households per month of which around one third are re-interviewed after 6 months. The time period that includes precise income information (previously income was only surveyed as bins) is July 1986 - December 2013. Overall, there are observations on 153,241 households (with or without re-interviews). We restrict the sample in the following way: (a) We only select households where the respondent is at most 65 years of age (excludes 30,701 observations). (b) We exclude observations with missing information on demographics (7,605 observations). (c) We exclude observations where the income is lower than the average unemployment benefits in that year (15,525 observations). (d) For households with re-interview we exclude households where the respondent changes between interviews (as identified by the demographics such as gender, age, education marriage status and racial background, excludes 2,901 observations). Moreover, we exclude households where the number of adults changes between interviews (excludes 3,182 observations). This restriction is made since we are analyzing per adult income in the household, so that changes in the number of adults in the household will reflect changes in this measure of income that might not be anticipated by respondents when they are asked about their income growth expectations.

Overall, this leaves a sample of 88,017 households for which we have full information on demographics as well as inflation expectations (sample INF). 17,500 of these households are both first interviewed in the second half of a year and have a re-interview (sample H2RE). This is the sample for which we have information on realized income growth. Out of sample INF, 41,742 households also provide income expectations and are first interviewed in the first half a year (sample H1), 44,010 provide income expectations and are first interviewed in the second half a year (sample H2).

Figure 15 shows how the income information in our sample compares to the income information in the Panel Study of Income Dynamics (PSID). The PSID is a panel survey that has been running since 1968 which has been widely used to analyze income dynamics. Plot (a) shows that in the first part of the sample real per capita income in the Michigan Survey is slightly lower than in the PSID. Since the late 1990s, however, the levels of income in both surveys are very similar. Note that we are not using the levels of income in our analysis. Instead, individual income growth rates are the center of our investigation. Plot (b) displays the distribution of these growth rates in the Michigan Survey and in the PSID. The distribution of income growth is very similar in both surveys.

A.2 Details about the Imputation Procedure

To increase the overlap of expectations and realizations we impute income growth realizations using the information of households with similar household characteristics who report their income growth for the relevant period. In the example of figure 1, all households interviewed for the first time between July 2002 and December 2002 report both their income in 2001 as well as their income in 2002. We can hence use their income in 2001 as well as all available household characteristics to predict their income growth 2001-2002. We then use this relationship to impute income growth 2001-2002 for all households interviewed for the

Figure 15: Comparison with Income Panel Study of Income Dynamics

(a) percentiles (b) distribution of income growth

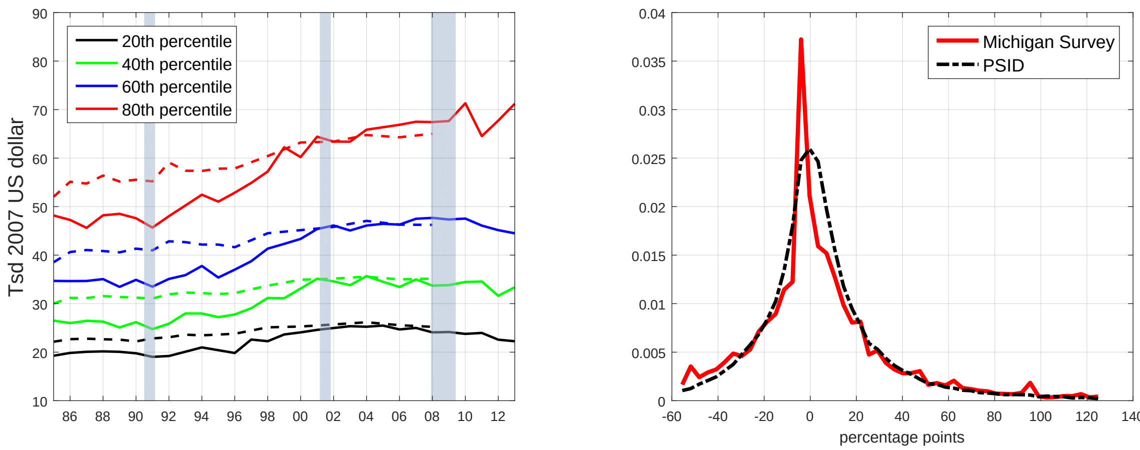

Note: The figure plots a comparison of reported income in the Michigan Survey and in the Panel Study of Income Dynamics (PSID). Plot (a) shows the percentiles of per capita real income over time: solid lines refer to the Michigan Survey distribution of income, dashed lines to the corresponding percentiles in the PSID. Plot (b) shows the distribution of real income growth rates in the Michigan Survey and in the PSID. Since the PSID changed to biannual surveys in 1997, the income growth rates have been constructed from PSID data 1986-1996 only.

first time in January 2002 to June 2002. The equation that we use to impute income growth realizations is the following:

$$g_{i,t+1} = \alpha + \beta X_{i,t} + \varepsilon_{i,t} \tag{19}$$

where $g_{i,t+1}$ is the growth rate in income of individual i from year t to year $t+1$ and $X_{i,t}$ includes a quadratic term in $\log(income_{i,t})$, a quadratic term in age, as well as indicators for education, gender, ethnic background, marriage status, number of adults, region, income growth expectations, inflation expectations and household weight in the survey. The imputation procedure is implemented as a multiple imputations algorithm using the predictive mean matching method with 5 nearest neighbors and 25 imputations. The imputation procedure is done separately for each survey year, using the observations from sample H2RE which report income changes for the respective year.

Figure 1(c) shows that for January households the overlap between expectation and imputed realization is now perfect. For February to June this overlap decreases but is still larger than the maximum overlap we obtain for July to December households on directly reported data. Moreover, for January to June households we do not need any re-interview so that we can use all observations in the data, not only the ones with re-interview. This greatly increases the sample size: We are able to obtain income growth realizations (and thus forecast errors) for the whole sample H1.

Furthermore, we can also increase the overlap for July to December households by imputing income changes for the following year. In the example of figure 1 we use the information provided by households interviewed for the first time in July to December 2003 to impute

income growth 2002-2003 for the households first interviewed in July to December 2002. This increases the overlap between their expectations and imputed realizations. The largest overlap is 11 months for December households, which is close to perfect. Note that for this step we base the imputation on the income that households reported in their second interview. Unlike in the case of the sample H1, we are hence only able to impute income changes for households who have a re-interview. Combined with the imputed sample H1 this generates the main sample of forecast errors of 58,369 observations (sample MAIN). Table 7 shows the distribution of imputed individual income growth rates in this sample compared to the directly reported income growth rates in sample H2RE. The distribution in the imputed data is very close to the distribution of the original data.

Table 7: Distribution of real reported income changes and imputed values

	mean	p5	p25	p50	p75	p95
directly reported	0.034	-0.378	-0.097	-0.015	0.133	0.572
imputed	0.032	-0.365	-0.103	-0.016	0.130	0.577

Note: The table compares the distribution of imputed individual growth rates in real income in sample MAIN with the growth rates in directly reported income in sample H2RE.

The main analyses reported in this paper are conducted on the sample MAIN where realized income growth has been imputed to maximize both the timing overlap and the number of observations. However, we have conducted robustness checks on the following subsamples: JAN (households with interview in January, income growth imputed, overlap perfect: 6,973 observations); DEC (households with first interview in December, income growth imputed, overlap close to perfect: 2,723 observations); JULY (households with interview in July, directly reported income growth, maximum overlap for directly reported data: 2,805 observations). The fact that the results hold on sample JULY where only directly reported data was used (no imputation) confirms that the results we find are not driven by the imputation procedure.

A.3 Time Series Plots of Errors in Nominal Income

Figure 16: Expectation errors in nominal income growth

(a) mean (b) mean by income quintile

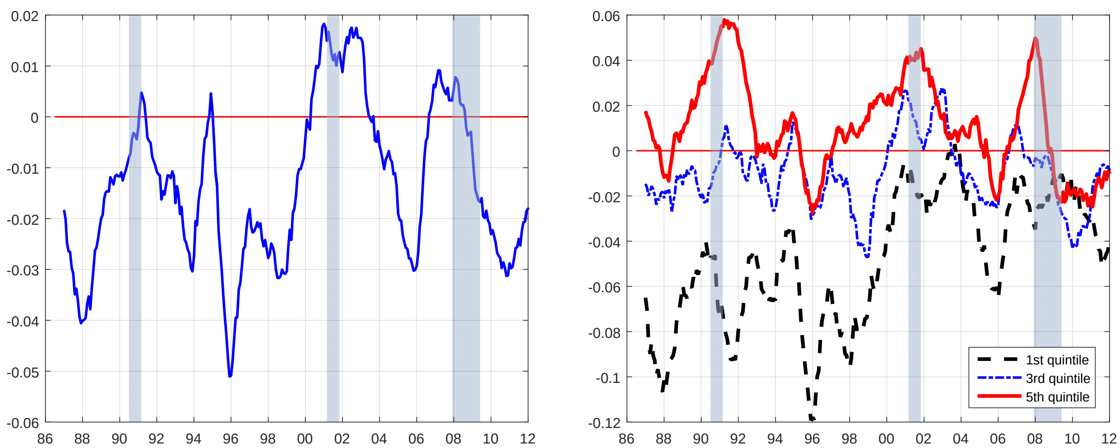

Note: The figure plots the 12-month moving average of mean expectation errors in individual nominal income growth. Expectation errors are winsorized at 5% and 95%. Data from the Michigan Survey of Consumers and own calculations. Grey areas represent NBER recessions. On the y-axis, 0.01 corresponds to 1 percentage point.

A.4 Interaction of Income with Age and Education

Table 8: OLS of forecast error on observables, interaction with education and age

	real	real	nominal	nominal
1st	−0.051***	−0.057***	−0.047***	−0.054***
	(0.007)	(0.010)	(0.007)	(0.010)
2nd	−0.017***	−0.021**	−0.016***	−0.018*
	(0.006)	(0.010)	(0.006)	(0.010)
4th	0.019***	0.027***	0.017***	0.025***
	(0.005)	(0.009)	(0.005)	(0.009)
5th	0.035***	0.047***	0.032***	0.043***
	(0.006)	(0.010)	(0.006)	(0.011)
no high school	0.013	0.023	0.019	0.030
	(0.014)	(0.027)	(0.014)	(0.028)
college	−0.014***	−0.008	−0.017***	−0.010
	(0.004)	(0.008)	(0.004)	(0.008)
age < 35	0.026***	0.021**	0.026***	0.021**
	(0.005)	(0.010)	(0.005)	(0.010)
50 ≤ age < 65	−0.013***	−0.015	−0.014***	−0.015
	(0.004)	(0.009)	(0.004)	(0.009)
1st × no high school		−0.019		−0.021
		(0.030)		(0.030)
2nd × no high school		−0.008		−0.011
		(0.034)		(0.035)
4th × no high school		0.015		0.013
		(0.037)		(0.038)
5th × no high school		0.020		0.021
		(0.045)		(0.046)
1st × college		0.005		0.003
		(0.013)		(0.013)
2nd × college		0.001		−0.000
		(0.012)		(0.013)
4th × college		−0.013		−0.011
		(0.011)		(0.011)
5th × college		−0.021*		−0.021*
		(0.012)		(0.012)
1st × age < 35		0.012		0.014
		(0.015)		(0.015)
2nd × age < 35		0.007		0.007
		(0.014)		(0.014)
4th × age < 35		−0.004		−0.005
		(0.012)		(0.012)
5th × age < 35		0.007		0.008
		(0.013)		(0.014)
1st × 50 ≤ age < 65		0.010		0.010
		(0.015)		(0.015)
2nd × 50 ≤ age < 65		0.005		0.003
		(0.014)		(0.014)
4th × 50 ≤ age < 65		−0.003		−0.004
		(0.012)		(0.012)
5th × 50 ≤ age < 65		−0.001		−0.001
		(0.012)		(0.012)
Month dummies	57498	57498	57498	57498

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Note: The table shows the results from the multiple imputations OLS regression of equation (2) (dependent variable is error in either real or nominal income growth on the household level) with additional interaction terms of income quintiles with education and age groups. Additional regressors (coefficients not shown) are a constant, racial background, number of adults in the household, gender, marriage status as well as region and month dummies.

A.5 Regression tables for error in aggregate unemployment expectation

Table 9: Ordered Logit / Ordered Probit of Unemployment Expectations

	(1) ologit	(2) oprobit
Income Quintile		
1st	−0.086***	−0.046***
	(0.023)	(0.013)
2nd	−0.032	−0.018
	(0.022)	(0.012)
4th	0.064***	0.036***
	(0.021)	(0.012)
5th	0.119***	0.069***
	(0.022)	(0.012)
Education		
no high school	−0.042	−0.017
	(0.039)	(0.022)
college	0.084***	0.048***
	(0.015)	(0.008)
Age		
age	−0.054***	−0.031***
	(0.005)	(0.003)
age × age	0.001***	0.000***
	(0.000)	(0.000)
Racial background		
black	−0.160***	−0.074***
	(0.029)	(0.016)
hispanic	0.078**	0.051***
	(0.035)	(0.020)
Number of adults		
1	−0.050	−0.025
	(0.030)	(0.017)
3 or more	0.083***	0.048***
	(0.024)	(0.014)
Other family characteristics		
female	−0.133***	−0.084***
	(0.014)	(0.008)
not married	−0.038	−0.024
	(0.028)	(0.016)
Region		
North Central	0.002	−0.002
	(0.020)	(0.011)
Northeast	−0.074***	−0.041***
	(0.022)	(0.012)
South	0.042**	0.023**
	(0.019)	(0.011)
Month dummies	yes	yes
Observations	96332	96332

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Note: The table shows the results from the ordered logit and ordered probit regression of categorical errors in individual expectations about aggregate unemployment development. The ordered categories are as follows: -2: far too pessimistic, -1: too pessimistic, 0: correct expectation, +1: too optimistic, +2: far too optimistic.

A.6 Regression tables for actual income changes and expected income growth

Table 10: OLS of actual income growth on observables

	(1) actual growth (real)	(2) actual growth (real)
Income Quintile		
1st		0.124***
		(0.011)
2nd		0.052***
		(0.009)
4th		−0.044***
		(0.007)
5th		−0.086***
		(0.009)
Education		
no high school	−0.027	−0.065***
	(0.016)	(0.017)
college	0.029***	0.074***
	(0.006)	(0.007)
Age		
age	−0.001	0.007***
	(0.002)	(0.002)
age × age	−0.000	−0.000***
	(0.000)	(0.000)
Racial background		
black	−0.041***	−0.052***
	(0.011)	(0.011)
hispanic	−0.021	−0.034***
	(0.013)	(0.013)
Number of adults		
1	0.033	0.077***
	(0.020)	(0.020)
3 or more	−0.013	−0.050***
	(0.010)	(0.010)
Other family characteristics		
female	−0.015**	−0.024***
	(0.006)	(0.006)
not married	−0.032*	−0.066***
	(0.018)	(0.018)
Region		
North Central	0.009	0.001
	(0.009)	(0.009)
Northeast	0.013	0.013
	(0.010)	(0.010)
South	0.010	0.005
	(0.009)	(0.009)
Constant	0.099**	−0.082*
	(0.044)	(0.045)
Observations	18181	18181
R^2	0.010	0.039

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Note: The table shows the results from the OLS regression of equation (2) where the dependent variable is actual real or nominal income growth on the household level. Estimation performed on sample of households that have two interviews with the first interview in the second half of the year.

Table 11: OLS of growth expectations on observables

	(1) expectation (real)	(2) expectation (real)	(3) expectation (nominal)	(4) expectation (nominal)
Income Quintile				
1st		0.017***		0.022***
		(0.002)		(0.002)
2nd		0.006***		0.009***
		(0.002)		(0.002)
4th		−0.001		−0.003**
		(0.002)		(0.002)
5th		0.003		−0.001
		(0.002)		(0.002)
Education				
no high school	−0.020***	−0.023***	−0.014***	−0.019***
	(0.003)	(0.003)	(0.003)	(0.003)
college	0.019***	0.022***	0.014***	0.019***
	(0.001)	(0.001)	(0.001)	(0.001)
Age				
age	−0.003***	−0.003***	−0.004***	−0.003***
	(0.000)	(0.000)	(0.000)	(0.000)
age × age	0.000***	0.000	0.000***	0.000
	(0.000)	(0.000)	(0.000)	(0.000)
Racial background				
black	0.011***	0.011***	0.017***	0.016***
	(0.002)	(0.002)	(0.002)	(0.002)
hispanic	−0.003	−0.005	0.004	0.002
	(0.003)	(0.003)	(0.003)	(0.003)
Number of adults				
1	−0.002	0.001	−0.002	0.003
	(0.003)	(0.003)	(0.003)	(0.003)
3 or more	0.006***	0.003	0.005***	0.001
	(0.002)	(0.002)	(0.002)	(0.002)
Other family characteristics				
female	−0.019***	−0.020***	−0.012***	−0.013***
	(0.001)	(0.001)	(0.001)	(0.001)
not married	0.013***	0.010***	0.014***	0.009***
	(0.003)	(0.003)	(0.003)	(0.003)
Region				
North Central	−0.018***	−0.018***	−0.018***	−0.018***
	(0.002)	(0.002)	(0.002)	(0.002)
Northeast	−0.011***	−0.011***	−0.011***	−0.011***
	(0.002)	(0.002)	(0.002)	(0.002)
South	−0.009***	−0.009***	−0.007***	−0.008***
	(0.002)	(0.002)	(0.002)	(0.002)
Constant	0.142***	0.125***	0.184***	0.160***
	(0.013)	(0.013)	(0.013)	(0.013)
Observations	89079	89079	93764	93764
R^2	0.045	0.046	0.044	0.047

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Note: The table shows the results from the OLS regression of equation (2) where the dependent variable is expected growth in real or nominal income growth on the household level. Estimation performed on full sample of households (with or without re-interview (first interview if there are two interviews), all interview months).

B Proof of Theorem 1

Income (net of age effects and the effects of other demographics) follows the process

$$Y_{it} = P_{it} \cdot V_{it} \tag{20}$$

$$P_{it} = P_{it-1}^{\rho} \cdot N_{it} \tag{21}$$

where P_{it} is a persistent component and V_{it} a transitory shock. Persistent income depends on past persistent income and a persistent shock N_{it}. Both shocks are independently and log-normally distributed with mean 1.

We assume that $1 > \hat{\rho} = \rho + \varepsilon > \rho$, so that all relevant moments exist and are finite. Expected income next period in this case is equal to $\mathsf{E}[Y_{it+1}] = \mathsf{E}[P_{it+1} \cdot V_{it+1}] = \mathsf{E}[P_{it}^{\hat{\rho}} \cdot N_{it+1} \cdot V_{it+1}] = P_{it}^{\hat{\rho}}$. Therefore the expected growth rate in income is $\mathsf{E}\left[\frac{\Delta Y_{it+1}}{Y_{it}}\right] = \frac{P_{it}^{\hat{\rho}} - Y_{it}}{Y_{it}}$ and the actual growth rate is equal to $\frac{\Delta Y_{it+1}}{Y_{it}} = \frac{P_{it}^{\rho} \cdot N_{it+1} \cdot V_{it+1} - Y_{it}}{Y_{it}}$ The expectation error can hence be calculated as:

$$\begin{aligned}
\psi_{it} &= E\left[\frac{\Delta Y_{it+1}}{Y_{it}}\right] - \frac{\Delta Y_{it+1}}{Y_{it}} \\
&= \frac{P_{it}^{\hat{\rho}} - Y_{it}}{Y_{it}} - \frac{P_{it}^{\rho} \cdot N_{it+1} \cdot V_{it+1} - Y_{it}}{Y_{it}} = \frac{P_{it}^{\hat{\rho}} - P_{it}^{\rho} \cdot N_{it+1} \cdot V_{it+1}}{Y_{it}} \\
&= \frac{P_{it}^{\rho+\varepsilon} - P_{it}^{\rho} \cdot N_{it+1} \cdot V_{it+1}}{Y_{it}} = \frac{P_{it}^{\rho}}{Y_{it}} \left(P_{it}^{\varepsilon} - N_{it+1} V_{it+1}\right) \\
&= \frac{P_{it}^{\rho-1}}{V_{it}} \left(P_{it}^{\varepsilon} - N_{it+1} V_{it+1}\right)
\end{aligned} \tag{22}$$

The *average* expectation error is then equal to $\mathsf{E}[\psi_{it}] = \frac{P_{it}^{\rho-1}}{V_{it}}\left[P_{it}^{\varepsilon} - 1\right]$. P_{it} can be re-written as a combination of its mean of $\mathsf{E}P = 1 + \bar{P}$ and the deviation from the mean p_{it}: $P_{it} = 1 + \bar{P} + p_{it}$. The term $\bar{P}$ is the lognormal mean correction term.

Using this notation, the expected error becomes

$$\mathsf{E}[\psi_{it}] = \frac{\left(1 + \bar{P} + p_{it}\right)^{\rho-1}}{V_{it}} \left[\left(1 + \bar{P} + p_{it}\right)^{\varepsilon} - 1\right] \tag{23}$$

For big enough current P_{it} (namely $p_{it} > -\bar{P}$), the term in the brackets is positive. This means that agents with income above this threshold on average overpredict their future income growth.

How does the expected error change with current P_{it}? $\frac{\partial \mathsf{E}\psi_{it}}{\partial P_{it}}$ has the same sign as $\frac{\partial F(z)}{\partial z}$ where $F(z) = z^{\rho+\varepsilon-1} - z^{\rho-1}$. We have

$$\begin{aligned}
F(z)' &= z^{\rho-1}[(\rho + \varepsilon - 1)z^{\varepsilon} - (\rho - 1)] \\
&\approx (\rho + \varepsilon - 1)\left[z^{\varepsilon} - \frac{\rho - 1}{\rho - 1 + \varepsilon}\right] \\
&= -(1 - \rho - \varepsilon)\left[z^{\varepsilon} - \frac{1 - \rho}{1 - \rho - \varepsilon}\right]
\end{aligned} \tag{24}$$

This expression is *positive* as long as $z^{\varepsilon} < \frac{1-\rho}{1-\rho-\varepsilon}$, that is as long as $z < \left(\frac{1-\rho}{1-\rho-\varepsilon}\right)^{1/\varepsilon}$. Because $\rho \gg \varepsilon$ and ε is close to zero, the expectation error is increasing in P_{it} until very very large values of current P_{it}. In the model calibration, we have $\rho = 0.9774$, $\varepsilon = 0.0057$, which translates into a threshold of $z \approx 1.4e22$.

C Imperfect information and forecasting shocks

Consider a simple state space model

$$x_t = \rho x_{t-1} + \eta_t \tag{25}$$
$$y_t = x_t + \mu_t \tag{26}$$

where η and μ are iid zero mean normal shocks with known finite variances. The forecasting error conditional on being in a particular quantile Q is $\mathsf{E}\big[y_{t+1} - y_{t+1|t} | y_t \in Q\big]$.

Suppose that t periods ago, the true state x_0 was known. It is then possible to write y_{t+1} as a function of starting state x_0, all previous η's and μ_{t+1}:

$$y_{t+1} = \eta_{t+1} + \mu_{t+1} + \rho\eta_t + \ldots + \rho^{t-1}\eta_1 + \rho^t x_0 \tag{27}$$

Similarly, $y_{t+1|t}$ can be written as a similar sum. However, now μ's also play role because of the imperfect information. It can be shown that

$$y_{t+1} - y_{t+1|t} = \eta_{t+1} + \mu_{t+1} + \rho\left[(1-K)\eta_t + K\mu_t\right] \tag{28}$$

where we assumed that the kalman gain K is does not change over time.[21] The conditional forecasting error behaves similar to

$$\mathsf{E}\left[y_{t+1} - y_{t+1|t} | y_t\right] \approx \mathsf{E}\left[\eta_t - \mu_t \middle| \eta_t + \mu_t + \sum_{\tau=1}^{t-1} \rho^{t+1-\tau}\eta_\tau\right] \tag{29}$$

However, $\eta_t - \mu_t$ is independent of $\eta_t + \mu_t$ and because the shocks are not serially correlated, $\sum_{\tau=1}^{t-1} \rho^{t+1-\tau}\eta_\tau$ does not overturn the fact that the term in the expectations is independent of the condition. Hence the conditional forecasting error is equal to the unconditional, which is equal to zero.

[21] This approximation is better the bigger t is at exponential rate.

D Simplified Setting: Model without Durables

The results from the main model also hold in a setting without durable goods. In this case the household optimization problem can be summarized as follows:

$$\max_{\{c_t\}_{t=0}^{\infty},\{s_t\}_{t=0}^{\infty}} \quad \sum_{t=0}^{\infty} \beta^t \, \mathsf{E}\big[U(c_t)\big] \tag{30}$$

$$\text{s.t.} \qquad c_t + s_t \leq R(s_{t-1}) + Y_t \tag{31}$$

$$R(s_t) = [1 + r\,(s_t)]s_t \tag{32}$$

$$r(s_t) = \begin{cases} r^l & \text{if } \; s_t > 0 \\ r^b & \text{if } \; -\kappa_y P_t \leq s_t \leq 0 \end{cases} \tag{33}$$

$$U(c) = \frac{c^{1-\gamma}}{1-\gamma}. \tag{34}$$

Income Y_t has the same functional form as in the main setting (equations (11) - (14)). Moreover, expectation biases are also modeled in the same way as in the full model (equations (15) and (16)). We keep the same parameter values as in the main setting.

Figure 17: Distribution of liquid assets in model without durables

(a) aggregate (b) lowest income (c) highest income

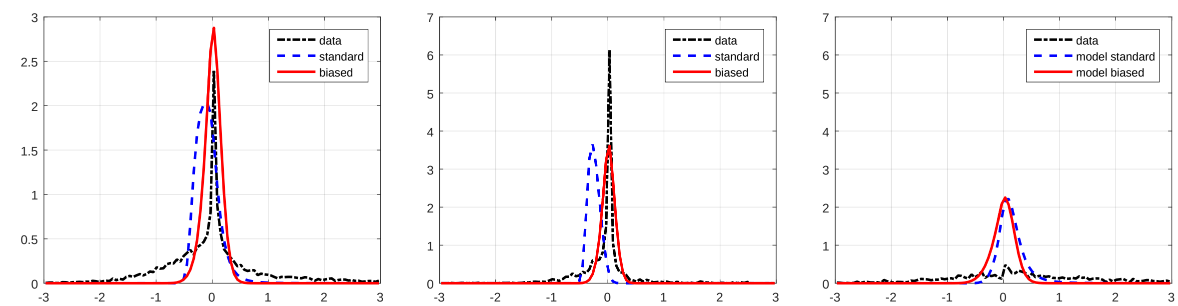

Note: The figure depicts the distribution of liquid assets in the simplified model versus data for the population as a whole and for different income quintiles. The panels show the data distribution (dashed black line) against the model distribution when households have non-rational expectations (solid red line). For comparison, the distribution under rational expectations is also plotted (solid blue line).

The first main implication of biased expectations in the main setting is that they allow the model to better fit the low end of the asset distribution. Figure 17 shows that this result holds in the restricted setting without durable goods. As in the full model, allowing for non-rational expectations helps to fit the liquid asset distribution. Households with biased beliefs are less willing to borrow than their rational expectations counterparts even though they face the same borrowing constraint. However, without the incentive to purchase durable goods households neither accumulate as much savings nor do they borrow as much as in the data. This simplified model hence retains the main result but has a worse model fit.

The second main implication of biased income expectations is that they alter the marginal propensities to consume. Under rational expectations, low income households have a MPC

Figure 18: MPC out of unexpected transfer in model without durables

(a) aggregate (b) lowest income (c) highest income

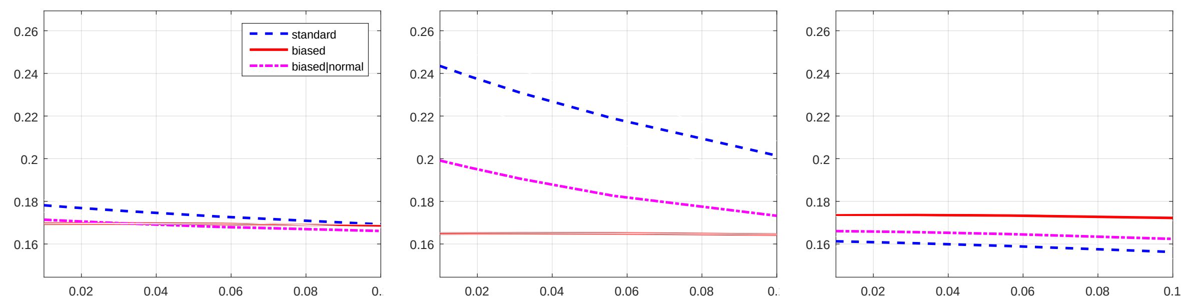

Note: The figure depicts the fraction of an unanticipated one-time transfer payment of varying sizes that is spent on non-durable consumption under different expectation scenarios: the red line depicts the MPC under biased expectations, the dashed blue line depicts the MPC under rational expectations. Panel (a) shows the MPC in the aggregate population while panels (b) and (c) show the MPC for the lowest and highest income quintile. Transfer sizes are expressed as fractions of average quarterly income in the economy.

which is much larger than the MPC of high income households. In contrast, this ratio is smaller if people have biased income expectations. Figure 18 shows that this result also holds in the restricted setting without durables.

E Results of calibration matching the Fully Rational Expectation Model

In this section, we choose the parameters to maximize the fit of the model with rational expectations and show that the results described in the main text still hold.

The only parameters which are free to be chosen differently compared to the benchmark model are the five parameters affecting the household preferences. The parameters describing the environment remain the same and the belief parameters are by assumption equal to the true process parameters.

Table 12: Parameter Values

Parameter		Value
preferences:		
discount factor	β	0.9875
risk aversion	γ	2
weight of durable goods in utility	θ	0.075
elasticity of substitution in utility	ξ	2.5
free durable services	$\bar{d}$	0.5

The resulting parameters are captured in table 12. Compared to the parametrization of the benchmark model with biased expectations, three parameters are different: the agents are more patient and more risk averse and there is less elasticity of substitution between durables and non-durable consumption. Figure 19 shows how well the fully rational model (and the corresponding version with biased expectations) fits the data. Table 13 summarizes the fit at selective quantiles.

Higher risk aversion combined with more patience makes the agents in the aggregate save more compared to the benchmark model. While the fit for the population as a whole is good, this specification remains to have counterfactual implications for liquid assets of low income households. Figure 20 shows that the model without rational expectations still generates too much borrowing for low income households, even though the preferences now lead to higher savings in the aggregate. Furthermore, the observation that the standard model generates much higher dispersion between the MPC of low and high income households also holds (see figure 21). The results described in the main text hence also go through when we allow the calibration to best fit the fully rational model to the data.

Table 13: Model fit

		quantile							mode
		0.05	0.10	0.25	0.50	0.75	0.80	0.90	
liquid assets	data	-1.29	-0.88	-0.30	0.03	0.76	1.36	5.46	-0.02
	model	-0.69	-0.52	-0.23	0.02	0.36	0.50	1.00	0.13
durables	data	0.13	0.20	0.39	0.79	1.43	1.62	2.21	0.23
	model	0.17	0.25	0.43	0.75	1.23	1.38	1.85	0.50

Note: Selected moments generated by the fully rational model compared to SCF.

Figure 19: Model fit, standard model parametrisation

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

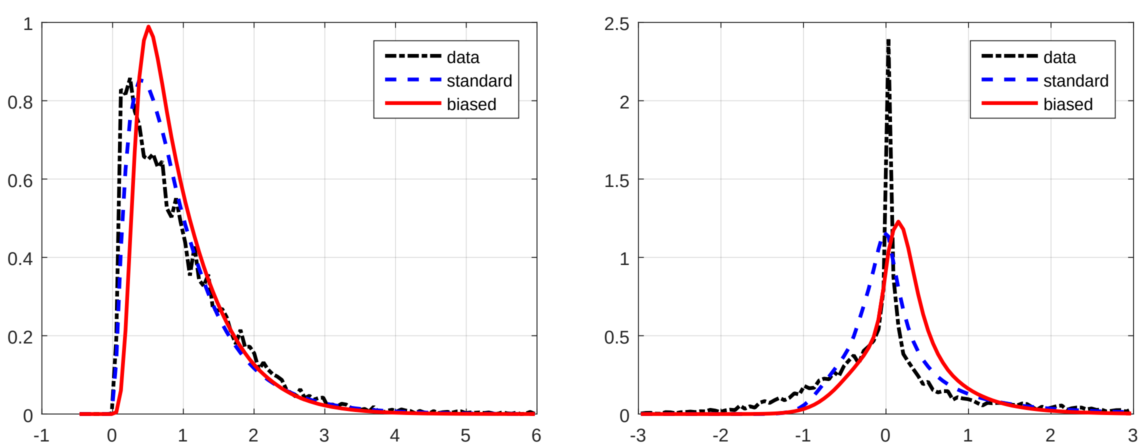

Note: The figure depicts the distribution for (a) durable goods and (b) liquid savings when the parameters are chosen to maximize the fit of the *standard* model. Data distributions (dashed black line) are compared to the distributions implied by model which allows for biased expectation (solid red line)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

Figure 20: Liquid assets by income quintile (s), standard model parametrisation

(a) 1st, rational versus both bias

(b) 5th, rational versus both bias

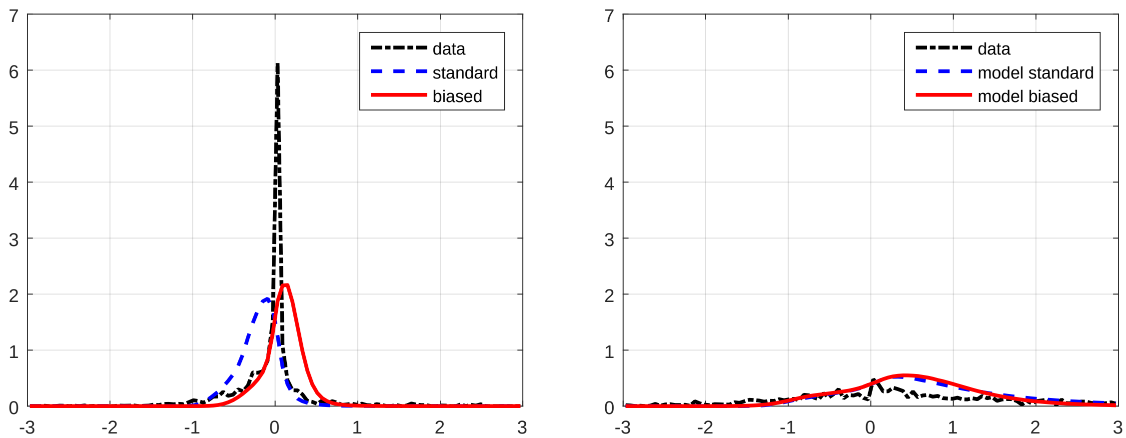

Note: The figure depicts the distribution of liquid assets in the model versus data for different income quintiles when the parameters are chosen to maximize the fit of the *standard* model. The panels show the data distribution (dash-dotted black line) against the model distribution when households have non-rational expectations (solid red line). For comparison, the distribution under rational expectations is also plotted (dashed blue line).

Figure 21: MPC out of unexpected transfer, standard model parametrisation

(a) aggregate

(b) lowest income

(c) highest income

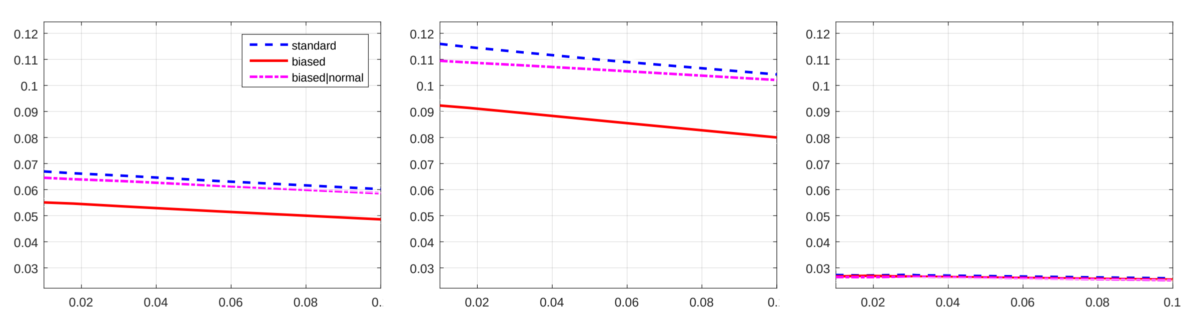

Note: The figure depicts the fraction of an unanticipated one-time transfer payment of varying sizes that is spent on non-durable consumption under different expectation scenarios: the red line depicts the MPC under biased beliefs, the dashed blue line depicts the MPC under rational expectations. Panel (a) shows the MPC in the aggregate population while panels (b) and (c) show the MPC for the lowest and highest income quintile. Transfer sizes are expressed as fractions of average quarterly income in the economy.

F Numerical implementation

F.1 Solution Algorithm

The model is solved using a value function iteration algorithm with Howard's Improvement. The solution of the rational agent's problem is standard. The policy functions of the agent with biased bieliefs are obtained in two steps. First, the problem is solved using the grid and transition matrices as if the biased beliefs were correct. After the solution converges, we do one more iteration of the value function iteration algorithm, now using the grid corresponding to the true data generating process, keeping the transition matrices and the continuation values $\mathsf{E}V'$ from the biased agent solution.

Including the discretization of the aggregate and idiosyncratic income components, we solve the model using the following grids:

- 90 grid points for liquid assets, the step size varies with the density (step size much smaller around zero)
- 60 grid points for durable assets, step size increasing with the level of durable asset.
- 15 states for the persistent idiosyncratic component P, levels and transition matrices generated using Rouwenhorst method
- 7 states for the idiosyncratic transitory component T, levels and probabilities generated using Gauss-Hermite Quadrature
- 2 states for the aggregate component Z, calibrated so the model delivers the same time spent in booms and recessions as the US economy.

Presence of the durable adjustment costs implies that the household has to decide whether to incur these costs and choose the optimal level of durable asset or let the durable good depreciate. In theory, in each step of the value function iteration, the values for both action and inaction have to be updated. Solving for the optimal action given adjustment is particularly costly, because it involves two-dimensional optimization. However, in practice it is not necessary to update both value functions at all grid points. If one keeps track of the boundary of the inaction region, both values only need to be updated in the neighborhood of the boundary. This step can lower the solution time considerably for well chosen grids, as the inaction region will occupy a large fraction of the state space.

F.2 Simulation

We obtain the distributions by simulating a panel of 125000 households for 1500 periods and discarding the first 500 observations. Using these 1000 periods, which include both booms and recessions as captured by the income component Z, we pool all the agents over all periods to construct the ergodic distributions.

To compute the marginal propensity to consume, we take all periods where the economy was in a recession, discard half and run the transfer experiment. We focus on the mean, because due to the non-linearity of the model, the actual consumption response is highly depending on other variables, both idiosyncratic (like the time since last durable purchase) and aggregate (the length of the recession) and hence vary for each individual in each recession.

F.3 Obtaining expectation errors in income growth

For all starting income levels $\{P,T\}_1$, we construct *all*[22] possible income path realizations and corresponding probabilities for 5 periods $\{P,T\}_1^5$. We then use this data in two ways. First, we use the first 4 periods of the income paths to construct income quintiles for the first year (the *annual* income in the first year Y^1 for a given realization is simply the sum of income in each quarter $Y^1 \equiv \sum_{t=1}^{4} P_t T_t$). This step also gives the probability distribution of $\{P,T\}_4$ conditional on being in a particular quintile of Y^1.

In the second step, we construct a second variable: annual income *conditional on the income state in the last quarter of the previous year*, Y^2. We do this by summing the income in periods two to five $Y^2 \equiv \sum_{t=2}^{5} P_t T_t$, remembering the corresponding probability and the starting state $\{P,T\}_1$. Again, we compute all possible values of Y^2 and the corresponding probabilities.

Finally, we combine the two pieces of information. We construct the growth rate as Y^2/Y^1, requiring that the last quarter of Y^1 is the same as $\{P,T\}_1$ used to compute Y^2. The expected growth rate conditional on being in a particular quintile is then computed by weighting all Y^2/Y^1 by the corresponding probabilities.

In order to find the belief parameters $\hat{\rho}$ and μ we proceed as follows: First we compute the expected growth rates for the true income process. Second, we iterate over guesses for $\hat{\rho}$ and μ until the implied expectation errors in growth rates correspond to the errors documented in the data.

[22] We discard any path with likelihood lower than $1e^{-9}$. The error in expectations introduced by this simplification is smaller than $1e^{-7}$.