

Household Consumption Forecast Errors and Self-Control Risk

Lars A. Lochstoer Stig R. H. Lundebj Zhaneta K. Tancheva *

July 1, 2026

Abstract

We show that households forecast future restraint that does not materialize. In the New York Fed Survey of Consumer Expectations, households more exposed to self-control problems underpredict consumption growth while overpredicting future improvements in personal finances; both errors widen during economic stress. Such opposite-signed errors are consistent with biased expectations about future self-control, but not with general pessimism or simple extrapolation from current conditions. We rationalize the joint forecast-error pattern in a model of state-dependent present bias and partial naiveté: households underestimate how much their future selves will overconsume and undersave, especially in bad states. In general equilibrium, the mechanism implies priced self-control risk even though agents forecast aggregate outcomes correctly.

JEL Codes: D91, E21, G12, G41

Keywords: Present bias, consumption expectations, self-control, discount-rate risk, household finance

*Lars A. Lochstoer: UCLA Anderson School of Management and NBER, lars.lochstoer@anderson.ucla.edu; Stig R. H. Lundebj: BI Norwegian Business School, stig.r.lundebj@bi.no; Zhaneta K. Tancheva: BI Norwegian Business School, zhaneta.k.tancheva@bi.no. This paper was previously circulated with the title “Present Bias and Discount Rate Risk.” We appreciate helpful comments from Harjoat Bhamra, Andrea Buffa, Julien Cujean, Bernard Dumas, Paul Ehling, Anastassia Fedyk, Mariana Khapko, Ali Lazrak, Ella Patelli, Ole Wilms, and participants at the WFA, SFS Cavalcade Asia-Pacific, FIRS, EFA, BFWG, BI-SHoF, Northwestern University, SAFE (Goethe University), Early Career Women in Finance Pre-WFA, SFA, Oslo Macro, and BCEA conferences, Adam Smith Workshop, BI Uncertainty Workshop, ESADE Spring Workshop, Finance and Accounting Symposium, and seminars at Boston University, BI Norwegian Business School, and Corvinus University. We acknowledge the support of the Centre for Advanced Study (CAS) in Oslo, Norway, which funded and hosted our research project (Asset Pricing with Heterogeneous Investors in Overlapping Generations) during the 2024/25 academic year. Any remaining errors are our own.

1 Introduction

Households do not only forecast aggregate conditions; they also forecast their own future behavior. These personal forecasts are economically important because many household choices require a prediction of one's future restraint: how much one will spend, how much one will save, whether one will repay debt, and whether one will follow through on financial plans. If households overestimate their future self-control, their consumption plans will be too restrained and their financial plans too optimistic.

We document precisely this pattern using the New York Fed Survey of Consumer Expectations (SCE). Households' realized consumption growth systematically exceeds their expected consumption growth. At the same time, realized personal financial conditions fall short of expected financial conditions. The two errors are opposite-signed for the same households: they expect to consume cautiously and to become financially better off, but subsequently consume more and improve their finances less than anticipated. This combination of low expected consumption growth and optimistic financial expectations is hard to reconcile with a uniform belief distortion, such as general pessimism, but is consistent with households believing that their future selves will be more disciplined than they turn out to be.

The empirical case rests on four facts. First, the underprediction of consumption growth is concentrated among households with *ex ante* characteristics plausibly associated with greater exposure to self-control problems: lower education, lower income, and weaker saving behavior. Quantitatively, the average household consumption-growth forecast error is about two percentage points per year, and the error is about four percentage points larger for the most self-control-exposed households than for the least exposed households. Second, the same households are too optimistic about their future personal financial situation. Third, both errors widen during economic stress. In bad times, consumption forecast errors for self-control-exposed households rise to roughly ten percentage points, while their financial expectations become even more optimistic relative to realizations. Fourth, in separate spending-expectations data, these households plan especially weak growth in non-essential relative to essential spending, and they are subsequently much more likely to experience financial distress.

We interpret the evidence through a model of state-dependent present bias and partial naiveté. Households have quasi-hyperbolic preferences and are partially naive about their own future self-control: they expect future selves to be more disciplined than they turn out to be. The model’s key object is the naiveté ratio, which summarizes the wedge between perceived and realized future self-control. The time variation in forecast errors disciplines the state dependence of this wedge. If the ratio is constant across states, the model can match the average forecast-error signs but not the good–bad swing. A state-dependent naiveté ratio can.

The quantitative household model has two roles. The one-asset model isolates the behavioral mechanism in the simplest environment. In the calibrated model, which already incorporates state-dependent income risk and rational beliefs about that risk, a state-independent naiveté ratio does not generate the cyclical forecast-error pattern. In the baseline calibration, bad states erode actual self-control while beliefs about future self-control do not adjust enough. This generates the empirical joint pattern: consumption forecast errors are positive and larger in bad states, while financial-condition errors are negative and more negative in bad states. The financial-condition pattern is not targeted; it follows because the same overconsumption that raises realized spending above planned spending lowers realized saving below planned saving.

The two-asset model asks whether the mechanism survives a more realistic household balance sheet. In a one-asset model, any precautionary buffer is liquid by construction. That is problematic for the relevant cohort, which is close to liquid-wealth constrained but has positive net worth held largely in illiquid form. We therefore add a liquid asset and an illiquid, higher-return asset with a withdrawal cost, following the hand-to-mouth structure in Kaplan and Violante (2014). The two-asset model preserves the forecast-error mechanism while matching the cohort balance sheet: liquid wealth is near zero and net worth lies in the empirical range. It also generates new testable implications. In bad states, naive households sharply cut contributions to commitment-like assets, while sophisticated households facing the same true deterioration in self-control sustain larger commitment inflows. Retirement-account contributions, mortgage prepayments, and other illiquid saving flows are natural empirical counterparts.

The paper also studies a stripped-down general-equilibrium implication. In a log-utility

endowment economy, agents correctly forecast aggregate quantities and asset payoffs but misforecast their own future behavior. Time-varying present bias changes the wealth-consumption ratio and creates self-control risk that is priced even when there is no aggregate consumption risk. Thus, when the errors in forecasting future self-control are state dependent and common across households, they can affect equilibrium discount rates. A rational benchmark with time-varying exponential discounting generates neither the same forecast errors nor the same self-control risk premium.

The paper contributes first to the literature on present-biased preferences, quasi-hyperbolic discounting, and self-control, beginning with Strotz (1956), Phelps and Pollak (1968), and Laibson (1997). A central prediction of this literature is that households may overconsume relative to prior plans unless they have commitment devices. Recent work studies the consumption, saving, monetary-policy, and welfare implications of these self-control problems in richer environments (Maxted, 2025; Maxted et al., 2025). We contribute direct survey evidence that self-control problems are visible in households' forecasts of their own consumption and financial conditions, and that the forecast errors are state dependent.

The paper also relates to work on partial naiveté and beliefs about future self-control. O'Donoghue and Rabin (1999) and O'Donoghue and Rabin (2001) show how naiveté about future present bias changes behavior, and recent work studies related implications for credit, contracts, and saving (DellaVigna and Malmendier, 2004; Heidhues and Köszegi, 2010; Eliaz and Spiegel, 2006; Groneck et al., 2024; Fedyk, 2025). We use household forecast errors to discipline the level and state dependence of the naiveté wedge. Work on imperfect foresight provides a complementary microfoundation: Gabaix and Laibson (2022) show how myopia can generate as-if hyperbolic discounting, with the intensity of the bias varying across states. The model is also consistent with evidence that stress, scarcity, and financial strain impair planning and cognitive control (Arnsten, 2009; Mani et al., 2013; Haushofer and Fehr, 2014; Sapolsky, 2017).

Finally, the paper speaks to asset-pricing models with time inconsistency and discount-rate risk. Luttmer and Mariotti (2003) show that constant present bias does not affect risk premia, while Khapko (2023) and Andries et al. (2024) study related effects of time-inconsistent

or horizon-dependent preferences. Our mechanism is different: agents can have rational expectations about aggregate outcomes and asset payoffs while making predictable errors about their own future consumption. The resulting discount-rate risk is therefore tied to household expectations about own behavior rather than to biased beliefs about aggregate cash flows.

The rest of the paper proceeds as follows. Section 2 presents the SCE evidence. Section 3 develops the household model and its calibration. Section 4 reports the one-asset and two-asset quantitative results. Section 5 discusses the two main comparisons: one asset versus two assets, and constant versus state-dependent naiveté. Section 6 presents the log-utility general-equilibrium benchmark, and Section 7 discusses implications for monetary-policy transmission. The final section concludes.

2 Household expectation evidence

This section provides the main empirical evidence. We document four facts consistent with self-control problems in household consumption expectations. First, households plausibly more exposed to self-control problems systematically underpredict their own total consumption growth. Second, the same households are simultaneously too *optimistic* about their future personal financial situation. Third, both errors widen during periods of economic stress: in bad times, self-control-exposed households underpredict their consumption growth even more, while becoming even more optimistic about their financial condition. Fourth, in separate spending-expectations data, the same households expect especially weak growth in non-essential relative to essential spending, and they are substantially more likely to subsequently experience financial distress. The juxtaposition of the consumption growth and financial situation forecast errors is central to our identification: generalized pessimism, negative sentiment, or worsening macroeconomic expectations would push both errors in the same direction, whereas we find errors of *opposite* sign that comove with economic stress. We close the section by showing that the pattern is not confined to financially constrained households and does not look like simple extrapolative expectations.

2.1 Data

We use data from the New York Fed’s Survey of Consumer Expectations, which is a survey of a population-representative rotating panel of 1,300 U.S. household heads, who own, buy or rent their home. The Core module includes monthly records of a total of 21,222 unique individuals over the sample period June 2013 until July 2023. Each consumer is observed for up to 12 consecutive months and there are about 1,000 respondents on average every month. The Spending module of the survey spans the period December 2014 until December 2022, but only takes place three times per year and includes 12,579 unique individuals.

We construct individuals’ realized real consumption growth and their expected real consumption growth to assess if the agents exhibit predictable forecast errors. The expected consumption growth is available in the Core module under the question: “By about what percentage do you expect your total household spending to [increase/decrease] over the next 12 months?”. In order to obtain expected real consumption growth we subtract expected inflation, which is also reported in the survey under the question: “What do you expect the rate of [inflation/deflation] to be over the next 12 months?”. The realized nominal consumption growth is recorded in the Spending module under the combined answers to the questions “How does your current monthly household spending compare with your household’s monthly spending 12 months ago?” and “In percentage terms, by how much has your current monthly household spending [increased/decreased] compared to 12 months ago?”. We obtain the realized real consumption growth after subtracting the actual inflation rate from FRED.

We complement the consumption growth forecast errors with an analogous forecast error for households’ own financial situation. The financial situation forecast error is defined as the difference between the reported personal financial situation over the past 12 months, based on the question “Do you think you are financially better or worse off than 12 months ago?”, and the predicted personal financial situation over the next 12 months, based on the question “Do you think you will be financially better or worse off 12 months from now?”. Negative values of this measure indicate that realized financial conditions fall short of prior expectations, i.e., that the household was overly optimistic about its future financial situation.

Table 1 reports the summary statistics of the main variables of interest from the raw, unfiltered New York Fed’s Survey data. In the first two rows we can already see that individuals’ average realized real consumption growth is higher than their expected real consumption growth (1.05% vs. -1.12%), which suggests that agents on average underestimate the amount they will consume next period. This is exactly the kind of expectation error we would expect from an agent that suffers from self-control problems. However, before conducting our formal statistical tests, we note that both variables have extreme outliers (e.g., maximum values above 9,000%, and minimum values below -100%). We therefore trim the data cross-sectionally at the 1% and 99% levels each time period, for which the respective values reported in Table 1 are reasonable. Finally, we limit the sample to individuals who respond to both the Core and the Spending modules to make sure we observe both their expected and realized consumption growth. After filtering the final sample includes 11,928 individuals.

The Core module also contains information about the individuals’ demographic characteristics. Education is recorded as the highest obtained degree ranging from Less than High School (1) to Professional Degree (8). We classify them into three categories, High School, Some College and College education, consistent with the classification of the New York Fed’s Survey, normalized to 1, 0.5 and 0, respectively. To ensure that the individuals have already completed their education at the time of the survey and to clean errors such as the maximum reported age of 891, we exclude individuals below the age of 25 and above or equal the age of 80. Afterwards, we classify individuals in age cohorts below 40, between 40 and 60, and above 60, consistent with the categories provided by the survey. Income ranges from Less than \$10,000 (1) to \$200,000 or more (11), that we classify in three cohorts consistent with the survey categories: Below \$50,000, Between 50,000 and \$100,000, and Above \$100,000, encoded as 1, 0.5 and 0, respectively. The risk tolerance is available in the Core survey under the question “On a scale from 1 to 7, how would you rate your willingness to take risks regarding financial matters?”. We group the individuals in three categories, low, medium and high risk tolerance: 1, 0.5 and 0, respectively. The last variable reported in Table 1 is an indicator of saving in general available in the Spending module under the question “People budget in different ways. Do you (and your family) generally

try to focus more on trying to save regular amounts of money?”. We define it such that it takes the value of 1, if individuals do not save in general, and 0 otherwise.

We also analyze how the individuals’ forecast errors vary with stress factors. In particular, we focus on state-level annual changes in unemployment rates available in FRED, as well as state-level economic conditions indicators. The unemployment rate is available on a monthly frequency at the beginning of the month, so we use the lagged value in order to make sure the information is available to the agents at the time of completing the survey. The economic indicators are developed by Baumeister et al. (2024) and available online at a weekly frequency. They are broad indices capturing the mobility, labor market, real activity, and financial conditions of households in the 50 U.S. states. The database does not include the District of Columbia. A value of zero indicates growth equal to the national long-run growth, negative values of the indicators correspond to lower than average growth and positive values denote higher than average growth. We multiply the variable by -1 such that positive values indicate lower than average growth and therefore higher degree of stress. We use the two-month lagged observations of the last week of the month in order to make sure the data is available at the time agents complete the survey.

2.2 Average forecast errors and self-control exposure

In this section we show that individuals systematically underpredict their real consumption growth and that the underprediction is concentrated among households more exposed to self-control problems.

The final sample gives us expected and realized annual real consumption growth for a broad set of individuals over the period from December 2014 to December 2022. To establish departures from rationality, we would ideally run regressions with the forecast error, $\Delta c_{t+1}^i - E_t^i(\Delta c_{t+1}^i)$, on the left hand side of the regression and predictors known at time t , x_t^i , on the right hand side. However, survey participants report their realized consumption growth over the last year along with the expectation of their consumption for the next year. This limits the specifications we are able to run to document the forecast error as we now have $\Delta c_t^i - E_t^i(\Delta c_{t+1}^i)$ on the left hand side of the regression, which can be predictable based on variables known at time t also under

rational expectations.

To overcome this challenge, we first establish that agents on average underpredict future consumption and that this underprediction is stronger for individuals that are likely to be more exposed to self-control problems. In particular, columns (1) and (2) of Table 2 give the estimate of the average forecast error μ from the regression

$$\Delta c_t^i - E_t^i(\Delta c_{t+1}^i) = \mu + \varepsilon_t^i. \quad (1)$$

The estimates are 1.6% and 2%, respectively, where the latter restricts the sample to only include the pre-Covid period. Thus, agents on average, across time and agents, expect their consumption growth to be substantially lower than the realized growth is. Since we only have 8 years of time series data available, one may worry that there are time trends in the data that the standard errors do not accurately capture. Columns (3) and (4) of the same table show regressions

$$\Delta c_t^i - E_t^i(\Delta c_{t+1}^i) = \mu_t + \mu_s + \beta^\top x_t^i + \varepsilon_t^i, \quad (2)$$

where μ_t and μ_s refer to time and state fixed effects, respectively, so the identification is cross-sectional, within-state. The x_t^i are agent i 's level of education, income, whether they save generally, age, and risk tolerance. The three former demographic variables are statistically significant at the 5% level or lower, while the two latter are not. All variables, except age, are coded into brackets between 0 and 1, as discussed in the previous section, where a higher value arguably indicates more self-control-exposed individuals. For instance, no high school education is coded as a 1, while a graduate degree is coded as a 0. The regression coefficients on the three significant coefficients are all positive, which indicates that more self-control-exposed agents underpredict consumption growth more, where the difference in forecast errors from low to high self-control exposure is about 3% per year for education, 2% per year for income and 0.5% per year for save generally.

The timing issue that we highlight above is not an issue for this regression as long as $Cov(\Delta c_t^i, x_t^i) = Cov(\Delta c_{t+1}^i, x_t^i)$, where the covariance is taken across time and agents, controlling

for time and state fixed effects. If the level of education, for instance, is set years before an agent enters the sample, this condition is likely to be satisfied. Similarly, the “Do you save regularly?” question refers to a longer-run condition of the agent that is unlikely to affect this covariance. The income brackets could be problematic, however, as an agent that had unexpectedly high income last year also might have consumed more than expected, but this would not necessarily affect the forward-looking expectation. In this case, the high-income agent would appear to underpredict consumption more than the low-income agent, which is the opposite of what we find. That said, the income brackets are very coarse, making it unlikely that an agent changes brackets so these shifts are in any case a small part of the variation in the data. Nevertheless, some agents likely do change income brackets (we cannot assess this directly with the data at hand as the survey participants only answer the income once), so this coefficient should be interpreted with caution. In the online appendix, we show that all our results are robust to removing the income variable from the regressions. In the online appendix, we also report analogous specifications using forecast errors constructed from the alternative survey waves available for each individual. Because respondents are observed predominantly over either two or three waves spanning four- or eight-month intervals, the expectation used to construct the forecast error is formed earlier than exactly 12 months before the realization. The results are qualitatively consistent with the main specifications.

2.3 Consumption growth and financial situation forecast errors

We now turn to the two central facts of the paper, presented jointly in Table 3: households more exposed to self-control problems underpredict their own consumption growth while simultaneously overpredicting improvements in their own financial situation, and both errors are amplified in periods of economic stress. Columns (1) and (2) of the table use consumption growth forecast errors and Columns (3) and (4) use financial situation forecast errors; Panel A reports unconditional estimates and Panel B reports the time variation in the errors with respect to state-level economic stress.

For these tests, we combine the three characteristics shown above to be significantly related

to underprediction – education, income, and saving generally – into a single composite measure of self-control exposure. The measure is computed as the predicted forecast error from a regression similar to that in Table 2 but using only these three variables (i.e., not age or risk tolerance); the construction is detailed in the online appendix. The measure is normalized such that agents with the highest measured self-control exposure are indicated by 1 and those with the lowest measured exposure by 0. We refer to this composite as the self-control-exposure cohort variable.

Unconditional errors. Panel A of Table 3 shows that the two forecast errors go in opposite directions. Column (1) repeats, for reference, the average consumption growth forecast error of 1.7% per year. Column (2) shows that, within state and time, moving from the least to the most self-control-exposed individuals raises the consumption growth forecast error by about 5.5 percentage points per year: more exposed individuals consume substantially more than they had predicted. In sharp contrast, Column (3) indicates that the average financial situation forecast error is significantly *negative*: on average, households’ realized financial situation falls short of what they had predicted a year earlier. Column (4) shows that this optimism is again concentrated among the self-control-exposed. Since the financial situation responses are categorical, we interpret the sign and relative magnitude rather than the units; the key observation is that the same households who underpredict how much they will spend simultaneously overpredict how well off they will be.

Time variation in errors. Panel B examines how both errors vary with economic conditions. Conditioning on stress variables makes the timing issue discussed in Section 2.2 more salient, as it entails conditioning on variables that affect the conditional distribution at time t of $t + 1$ outcomes. To address this, we create cohorts based on education, income, and saving generally, as well as the composite self-control-exposure variable. Importantly, these characteristics are not changing with the conditioning variables. We create cohorts by state and compute, at each time t , the average realized and expected consumption growth (and analogously the financial situation outcomes) for each cohort in each state, $\Delta c_{s,t}^j$ and $E_t^j(\Delta c_{s,t+1})$, where j refers to the cohort. We then run regressions of the form:

$$\Delta c_{s,t+1}^j - E_t^j(\Delta c_{s,t+1}) = \mu_t + \mu_s + \beta_1 x_t^j + \beta_2 z_{st} + \beta_3 x_t^j z_{st} + \varepsilon_{s,t+1}^j, \quad (3)$$

where x_t^j is the self-control-exposure cohort variable and z_{st} is a state-level indicator of economic stress – either the state-level economic conditions indicator (see Baumeister et al. (2024)) or the state-level change in unemployment. In the regressions, we sign both variables such that high values indicate bad times. Note that in this regression the timing of the left hand side variable is the usual definition of a forecast error due to the use of representative cohorts to construct the realized and expected outcomes. The same specification is run with the financial situation forecast error as the dependent variable in Columns (3) and (4).

The main coefficient of interest is β_3 , which asks whether the sensitivity of the forecast errors to self-control exposure changes with economic conditions. For consumption growth forecast errors, a positive β_3 implies that more self-control-exposed agents underpredict their future consumption growth more in bad times. This is indeed the case for both stress indicators (Columns (1) and (2)): the interaction coefficients are 0.90 and 0.78 and statistically significant. In terms of magnitudes, the most self-control-exposed agents have an average annual consumption growth forecast error of about 6%, relative to the least exposed agents who do not make any mistakes on average, and that error increases to about 10% if the stress indicators rise by 2 standard deviations. For the financial situation forecast errors (Columns (3) and (4)), the pattern is reversed: the interaction coefficients are negative and statistically significant, implying that self-control-exposed individuals become increasingly *optimistic* about their future financial condition precisely in bad times, even as their realized outcomes deteriorate.

This contrast is central for interpretation. Explanations based on general pessimism, negative sentiment, or worsening macroeconomic expectations would predict that both the consumption expectations and the financial outlook of exposed households deteriorate together in bad times. Instead, exposed households expect to consume less than they end up consuming while expecting to be better off than they end up being – a combination that points to misprediction of their own future behavior rather than misperception of their environment. Notably, the coefficients on the stress indicators themselves (β_2) are insignificant throughout, indicating that the forecast errors of the least exposed cohorts are not related to aggregate (in this case, state-level) economic outcomes in general.

The estimates in Panel B, Column (1) also discipline the calibration of the model in Section 3.7. The annualized average error of the most exposed cohort is the average prediction based on the significant explanatory variables from the regression model using the economic conditions indicator at a weekly frequency, the annualized standard deviation of the forecast errors is estimated as the projected standard deviation of the economic conditions indicator in the regression model, and the persistence of the forecast errors is based on the annualized autocorrelation of the economic conditions indicator. We use these values as an upper bound for the forecast errors we calibrate our model to.

In the online appendix, we show that the results are not driven by the use of trimming in consumption outcomes or the inclusion of the income variable in the composite measure. We also verify that the results are not driven by the Covid period: excluding the Covid year 2020 from the sample, both the unconditional forecast errors in Panel A and their time variation with economic stress in Panel B remain statistically significant and of similar magnitude. As an additional robustness check, the online appendix also presents analogous specifications using forecast errors constructed directly at the individual level from the available survey waves rather than cohort averages. Since respondents are typically observed over either four- or eight-month intervals, the corresponding expectations are formed less than 12 months before the realization date. The results remain qualitatively similar to the baseline cohort-based specification.

2.4 Mechanisms: discretionary spending and financial distress

The forecast-error evidence above shows that self-control-exposed households mispredict their own future outcomes in a particular, signed way. In this section we use additional SCE modules to examine the behavior underlying these errors. We show that the same households plan restraint precisely on the discretionary spending margin where temptation is most relevant, and that they subsequently experience financial distress at much higher rates – the footprint of overconsumption relative to plan.

Discretionary versus essential spending. We first analyze expectations for essential and non-essential consumption, constructed from survey questions available in the Spending module on

“everyday spending on essential items over the next 12 months... related to what you absolutely need” and “non-essential spending (such as on hobbies, leisure, vacation, and other items that you do not absolutely need)”. Panel A of Table 4 shows that individuals in the self-control-exposure cohort report significantly lower expected growth in non-essential consumption than the least exposed individuals, and that this gap is much larger for non-essential than for essential spending: the cohort coefficient is -10.4 for non-essential expectations, -3.7 for essential expectations, and -6.1 for the difference between the two. Table 5 further shows that this differential planned restraint intensifies in periods of economic stress: the interaction between the cohort variable and both stress indicators is negative and statistically significant, and again largest for non-essential consumption. Since non-essential spending is more discretionary and subject to temptation, this pattern is consistent with planned restraint on the margin where self-control problems are likely to be most relevant. In other words, these results are about expected spending growth rather than realized-minus-expected forecast errors: the self-control-exposure cohort expects to restrain non-essential spending relative to essential spending, and increasingly so in bad times. Borrowing constraints can restrict feasible consumption, but they do not by themselves predict this differential pattern in planned non-essential versus essential spending under rational expectations.

Ex post financial distress. Panel B of Table 4 shows that the same individuals are significantly more likely to experience financial distress, based on the questions “Over the past 12 months, did you max out (borrow up to the limit) on any of your credit cards?” and self-reported indicators of being late on loan payments available in the Credit Access module. Moving from the least to the most self-control-exposed individuals raises the probability of having maxed out a credit card by about 16 percentage points, of being more than 30 days late on a loan payment by about 10 percentage points, and of being more than 90 days late by about 31 percentage points.

Taken together with the forecast-error evidence, these results form a coherent picture of misprediction of own behavior. Self-control-exposed households plan restraint, especially on discretionary spending and especially in bad times; they expect their financial situation to improve; yet they end up consuming more than planned, their financial situation falls short of expecta-

tions, and they disproportionately end up in financial distress. Individuals consume more than they had anticipated in the aggregate and subsequently face tighter financial conditions – precisely the joint pattern implied by self-control problems combined with imperfect awareness of those problems.

2.5 Alternative explanations

2.5.1 Financial constraints

Financial constraints are an important confound, but not because they provide a rational-expectations explanation for predictable signed forecast errors. In a rational-expectations benchmark, constrained households may have lower consumption levels, higher marginal propensities to consume, or consumption that is more sensitive to income and spending shocks. If they correctly understand their constraints and the distribution of future shocks, however, their forecast errors should still be mean zero conditional on the information used to form the forecast. The concern is therefore that the empirical pattern might be mechanically driven by financially distressed or hand-to-mouth households with especially volatile consumption, rather than by self-control problems. To address this concern, we repeat the analysis in a subsample of financially unconstrained households.

Questions relevant for defining financial constraints come from the Credit Access module of the survey. In addition to our standard filters, we restrict the sample to respondents participating in this module (7,546 individuals). We classify individuals as financially constrained—and exclude them from the sample—if they satisfy any of the following conditions: (i) they answer “Yes” to “Over the past 12 months, did you max out (borrow up to the limit) on any of your credit cards?”; (ii) they report “I did not think I would get approved” as the reason for not applying for credit in the past 12 months; (iii) they give the same response when asked why they are unlikely to apply for credit in the next 12 months; or (iv) they report “No, my request was rejected” to whether a credit application was granted.

The results are presented in Table 6. First, more self-control-exposed financially unconstrained individuals make significant consumption growth forecast errors of about 6% per year,

similar to those in the full sample. Second, these forecast errors are time-varying and increase significantly in times of stress, based on the interaction term coefficient β_3 . These results are also robust to excluding income from the composite self-control-exposure cohort variable. The evidence suggests that the forecast-error pattern is not confined to financially constrained households. This distinction matters: borrowing constraints restrict feasible consumption and can amplify consumption responses to shocks, whereas our mechanism concerns households' miscalculation of their own future restraint.

2.5.2 Extrapolative expectations

A second natural alternative is that households form expectations by extrapolating from recent outcomes rather than miscalculating their own behavior. We argue that the joint pattern of our four facts is difficult to reconcile with extrapolation along the sign of the errors and their comovement with economic stress.

First, consider the unconditional evidence. Extrapolation from past outcomes generates forecast errors of both signs depending on the recent history: a household that extrapolates high recent consumption growth overpredicts future growth when growth mean-reverts, and symmetrically underpredicts after low recent growth. Averaged over a sample that spans both good and bad histories, extrapolation therefore produces forecast errors of both signs that largely offset, rather than a persistent one-signed error. In the data, realized consumption growth systematically exceeds expectations – the errors are predominantly positive throughout the sample.

Second, and most importantly, extrapolation cannot account for the opposite-signed comovement of the two forecast errors with economic stress. A household that extrapolates current conditions forward should become uniformly more pessimistic in bad times – about its consumption, its income, and its financial situation alike. Even a mean-reversion-based account, in which depressed current consumption generates positive consumption growth surprises in bad times, would imply that expectations about the financial situation err in the same pessimistic direction. Instead, we find the opposite: in bad times, self-control-exposed households underpredict their consumption growth even more while becoming even more *optimistic* about their future financial

condition, and they simultaneously plan greater restraint specifically on non-essential spending. No single belief-formation rule applied uniformly to all of a household’s expectations delivers errors of opposite sign across outcomes for the same household at the same time. Misprediction of own future behavior does: households plan restraint on the temptation margin, expect that restraint to restore their financial balance, and are then surprised on both fronts when the restraint does not materialize.

3 The model setup

3.1 Overview

We build a quantitative model of household consumption designed to explain the empirical results reported in the previous section: high-bias households jointly (i) *underpredict* their own future consumption and (ii) *overpredict* their own future financial condition, with both errors larger in bad aggregate states.

In the model, a single self-control wedge—households consume more than their forward-looking selves had planned—at once leads households to underpredict consumption and overpredict wealth, because the same overspending that pushes realized consumption above its forecast pushes realized saving below it. The natural rival, in which households simply hold pessimistic beliefs about their future income or financial environment, moves the two errors the *same* way: a household that expects to be poorer than it turns out to be under-predicts both its spending and its wealth. The opposite-signed pattern in the data—pessimism about one’s own consumption together with optimism about one’s own financial condition—is therefore the signature of a misperceived *self-control* problem rather than a misperceived environment.

The model’s key novel ingredient is that the gap between a household’s actual present bias and its belief about that bias varies with the aggregate state of the economy. In the baseline, bad states erode actual self-control while beliefs about future self-control do not adjust—households carry self-control risk they fail to anticipate. Everything else is taken from the literature: quasi-hyperbolic discounting (Laibson, 1997), partial naivete (O’Donoghue and Rabin, 1999, 2001),

the recursive partially-naive continuation-value construction of Maxted (2025) and Maxted et al. (2025), a countercyclical earnings process in the tradition of Guvenen et al. (2021), and, in the extension, the liquid/illiquid two-asset structure of Kaplan and Violante (2014).

3.2 Preferences and partial naivete

A household has per-period utility $u(c) = c^{1-\gamma}/(1-\gamma)$ with $\gamma > 0$, and discounts the future quasi-hyperbolically. We calibrate the model at a monthly frequency. The self that acts in period t evaluates streams according to

$$u(c_t) + \beta \mathbb{E}_t \sum_{j \geq 1} \delta^j u(c_{t+j}), \quad \beta \in (0, 1], \quad \delta \in (0, 1), \quad (4)$$

where β is the *true* present-bias parameter—indexed below by the aggregate state, β_z —and δ the standard (exponential) discount factor. Present bias makes the household time-inconsistent: the self at t would like later selves to save more than those selves will choose to.

Households are *partially naive* in the sense of O’Donoghue and Rabin (2001): the self at t correctly knows the long-run discount δ but believes its future selves will act with a present-bias parameter $\beta^E \in [\beta, 1]$, where $\beta^E = \beta$ is full sophistication and $\beta^E = 1$ is full naivete. Our single departure from the literature is to let the *naivete ratio* β_z^E/β_z —the gap between believed and realized future self-control—depend on the aggregate state $z \in \{G, B\}$ (good, bad). The forecast-error data identify the state dependence of this ratio. In the *baseline* the state dependence sits on the true side:

$$\beta_B < \beta_G, \quad \beta_z^E = \beta^E \text{ constant}, \quad (5)$$

so actual self-control deteriorates in bad states while beliefs about future self-control do not adjust. Households carry *self-control risk*: bad times make them more present-focused than they expect to be—the representation aligned with evidence that stress and scarcity impair forward-looking self-control (Mani et al., 2013; Haushofer and Fehr, 2014). Section 4 evaluates the two alternative representations of the same ratio path—a constant true bias with state-dependent beliefs (β fixed, $\beta_B^E > \beta_G^E$), and a constant sophistication share ($\beta_z^E = \lambda\beta_z + 1 - \lambda$)—and shows

that all three are nearly equivalent on the survey moments while differing sharply in psychology, feasibility, and implications for actual behavior.

3.3 Forecast behavior, actual behavior, and the forecast error

We use the recursive construction of Maxted (2025) and Maxted et al. (2025). Let x collect the household's individual states (assets and the idiosyncratic income state) and let z be the aggregate state. Define the *perceived continuation value* $\widehat{V}_z(x)$ as the value of the sophisticated Markov-perfect equilibrium that would be played by a sequence of selves each carrying present-bias β_z^E . The household believes it will follow this equilibrium; we therefore call its policy the *forecast policy* and denote it $\widehat{c}$. Beliefs are Markov-consistent: the household believes a future self in aggregate state z' will act with bias $\beta_{z'}^E$ (a constant β^E in the baseline). Its naivete is the gap between this belief and the bias that will actually govern behavior, $\beta_{z'}$ —in the baseline, a failure to anticipate that bad states will erode its own self-control. It solves, state by state,

$$\widehat{c} = \arg \max_c \left\{ u(c) + \beta_z^E \delta \mathbb{E}[\widehat{V}_{z'}(x') \mid x, z] \right\}, \quad (6)$$

and $\widehat{V}$ is the fixed point generated by $\widehat{c}$, $\widehat{V}_z(x) = u(\widehat{c}) + \delta \mathbb{E}[\widehat{V}_{z'}(x')]$.

The household's *actual* behavior, however, is governed by its true present bias. Each period the acting self takes a one-shot deviation: it applies the true bias in force, β_z , today while mistakenly continuing to value the future through the same perceived continuation value $\widehat{V}$. The actual policy c^* solves

$$c^* = \arg \max_c \left\{ u(c) + \beta_z \delta \mathbb{E}[\widehat{V}_{z'}(x') \mid x, z] \right\}. \quad (7)$$

The wedge between forecast and actual consumption is transparent in the homothetic interior benchmark, or locally when the two choices are evaluated at a common marginal continuation value. Let $\mathcal{M}_z(x)$ denote the marginal value of saving into the perceived continuation problem.

The forecast first-order condition is

$$u'(\widehat{c}) = \beta_z^E \delta \mathcal{M}_z(x),$$

whereas the actual first-order condition is

$$u'(c^\star) = \beta_z \delta \mathcal{M}_z(x).$$

Taking the ratio gives

$$\frac{u'(c^\star)}{u'(\widehat{c})} = \frac{\beta_z}{\beta_z^E}.$$

With CRRA utility, $u'(c) = c^{-\gamma}$, so

$$\left(\frac{c^\star}{\widehat{c}}\right)^{-\gamma} = \frac{\beta_z}{\beta_z^E},$$

or

$$\frac{c^\star}{\widehat{c}} = \left(\frac{\beta_z^E}{\beta_z}\right)^{1/\gamma} > 1. \tag{8}$$

Thus, partial naiveté implies that the household consumes more than it forecasts: the agent believes that future selves will be more disciplined, consume less, and save more. In the full stochastic model with income risk, borrowing wedges, and asset adjustment margins, $c^\star$ and $\widehat{c}$ generally imply different next-period states, so the marginal continuation value differs across the two choices and the proportionality in (8) need not hold globally. We therefore solve the forecast and actual policies directly from their own first-order conditions and use (8) only as a diagnostic benchmark.

The forecast-error. The SCE Household Spending Survey elicits the household’s expected *percent growth* in total household spending over the next twelve months (Federal Reserve Bank of New York, 2024). From a common initial condition (x_0, z_0) we simulate two twelve-month paths that face the *same* realized income shocks—an actual path on $c^\star$ and a forecast path on

$\hat{c}$ —and form annual spending

$$C_{t,t+12}^* = \sum_{m=1}^{12} c_{t+m}^*, \quad \hat{C}_{t,t+12} = \sum_{m=1}^{12} \hat{c}_{t+m}. \quad (9)$$

The survey object is forecasted growth in annual spending, $\Delta \log C_{t,t+12}$, relative to the already-realized prior year $C_{t-12,t}$. Because that baseline is common to the realized and the forecast growth rate, it cancels, and the household-level forecast error in annual-spending growth is

$$\text{FE}_z = \mathbb{E}[\Delta \log C_{t,t+12} - F_t \Delta \log C_{t,t+12} \mid z_0 = z] = \mathbb{E}[\log C_{t,t+12}^* - \log \hat{C}_{t,t+12} \mid z_0 = z], \quad (10)$$

a percentage object with cyclical swing $\text{FE}_B - \text{FE}_G$.

The financial-condition error is defined symmetrically on end-of-year liquid wealth, normalized by annual income,

$$\text{FC}_z = \mathbb{E}\left[\frac{b_{12}^* - \hat{b}_{12}}{Y_0^{\text{ann}}} \mid z_0 = z\right], \quad (11)$$

where Y_0^{ann} is annual permanent income. Since the household saves less than it forecasts, $\text{FC}_z < 0$: it over-predicts its financial condition.

3.4 The income process

Income is normalized by permanent income, in the manner standard since Carroll (1997). Log permanent income follows a random walk with monthly innovation ψ , so that normalized cash-on-hand inherits the term $1/\psi'$. Conditional on employment, transitory income is ξ ; an employment state $s \in \{E, U\}$ evolves as a two-state Markov chain whose hazards depend on the aggregate state, and unemployed households receive a replacement income b_{UI} . The aggregate state $z \in \{G, B\}$ follows a two-state chain Π calibrated to U.S. unemployment dynamics; its stationary probability of the bad state is 0.38. All innovations have larger dispersion and the unemployment hazards are higher in the bad state, generating countercyclical earnings risk with a fatter lower tail in recessions, consistent with Guvenen et al. (2021).

3.5 The one-asset model

In the one-asset economy the household holds a single asset that can be saved at gross return R_s or borrowed at $R_b > R_s$ up to a limit. Let m be cash-on-hand and a end-of-period assets. The budget and law of motion are

$$c + a = m, \quad a \geq \underline{a}, \quad m' = \frac{R(a)}{\psi'} a + y', \quad R(a) = \begin{cases} R_s, & a \geq 0 \\ R_b, & a < 0, \end{cases} \quad (12)$$

where y' is realized normalized income. The forecast and actual policies solve (6) and (7) with $x = (m, s)$.

3.6 The two-asset model

The one-asset model identifies the behavioral wedge and is all that we need to illustrate the model's main channel. The two-asset model shows that the same mechanism survives in a realistic household balance sheet and generates additional testable implications for liquid wealth, MPCs, and commitment-saving flows.

The two-asset economy adds an illiquid asset, following Kaplan and Violante (2014) but kept deliberately minimal: there is no fixed adjustment cost, no housing service flow, and no default. The household holds liquid balances m (saved at R_s , borrowed at R_b) and illiquid balances k that earn gross return $R_k > R_s$ but can be reduced only by paying a *proportional* withdrawal cost ϕ . Deposits into the illiquid account are free. Let a denote post-decision illiquid holdings. The within-period budget is

$$x = \begin{cases} m - (a - k), & a \geq k \quad (\text{deposit}) \\ m + (1 - \phi)(k - a), & a < k \quad (\text{withdraw}), \end{cases} \quad c + b = x, \quad (13)$$

where b is end-of-period liquid wealth, and the laws of motion are

$$m' = \frac{R(b)}{\psi'} b + y', \quad k' = \frac{R_k}{\psi'} a. \quad (14)$$

The forecast and actual policies solve (6) and (7) with $x = (m, k, s)$ and an additional choice of a . The illiquidity wedge $1 - \phi$ creates a region of inaction in which the household neither deposits nor withdraws; this region is the source of the numerical difficulty discussed in Section A.

3.7 Calibration

The calibration is hierarchical: each block is disciplined by a different class of evidence, and the behavioral parameters are calibrated last, to a target that the rest of the model does not touch.

(1) True present bias is anchored to the literature. The decision period is monthly, and the present-bias wedge β_z applies at each decision date (it is an immediate weight on the future, not compounded across months as δ is). We anchor its ergodic mean to the baseline value of the quasi-hyperbolic consumption literature, $\mathbb{E}_z[\beta_z] = (1 - \pi_B)\beta_G + \pi_B\beta_B = 0.70$ (Laibson, 1997; Maxted, 2025; Maxted et al., 2025), where $\pi_B = 0.38$ is the stationary bad-state probability of the aggregate chain. Conditional on all other calibrated parameters and on the model’s numerical mapping from policies to annual consumption-growth forecast errors, the three unknowns $(\beta^E, \beta_G, \beta_B)$ are pinned down by three moments: $\text{FE}_G = 0.044$, $\text{FE}_B = 0.100$, and $\mathbb{E}_z[\beta_z] = 0.70$. The state variation in true self-control is thus not externally measured; it is inferred from the forecast errors. The standard discount factor is set to $\delta = 0.96$ annually, so $\delta_m = \delta^{1/12}$, and we set risk aversion to $\gamma = 2$. We also report the analytically transparent $\gamma = 1$ case of Maxted (2025) as a benchmark.

(2) The income process. We calibrate the household’s income process to evidence in Guvenen et al. (2021). The income process has three components: permanent income, a transitory shock,

and employment status. Let the aggregate state be $z \in \{G, B\}$. Permanent income evolves as

$$P' = \psi' P, \quad \log \psi' \mid z' \sim \mathcal{N}\left(-\frac{1}{2}\sigma_{\psi,z'}^2, \sigma_{\psi,z'}^2\right),$$

so that $\mathbb{E}[\psi' \mid z'] = 1$. Employment status $s \in \{E, U\}$ follows a two-state Markov chain with transition probabilities that depend on the aggregate state:

$$\Pr(s' = U \mid s = E, z') = \pi_{z'}^{EU}, \quad \Pr(s' = E \mid s = U, z') = \pi_{z'}^{UE}.$$

Bad states have higher job-loss probabilities and lower job-finding probabilities.

Level labor income is $Y' = P'y'$, where normalized income is

$$y' = \begin{cases} \xi', & s' = E, \\ b_{UI}, & s' = U. \end{cases}$$

When employed,

$$\log \xi' \mid z' \sim \mathcal{N}\left(-\frac{1}{2}\sigma_{\xi,z'}^2, \sigma_{\xi,z'}^2\right),$$

so that $\mathbb{E}[\xi' \mid z'] = 1$. When unemployed, the household receives replacement income b_{UI} as a fraction of permanent income.

We solve the model in variables normalized by permanent income, following Carroll (1997). Since $P' = \psi' P$, dividing the level budget constraint by P' introduces the factor $1/\psi'$ multiplying beginning-of-period assets in the normalized laws of motion. In the baseline, the large downside income realization is recoverable unemployment: an unemployment spell reduces normalized income to b_{UI} , but permanent income continues to evolve independently, so income recovers upon re-employment. Section 4 considers robustness variants that change one element of this structure at a time: loading the downside onto permanent income, fattening the transitory-income tail, or raising b_{UI} .

The parameters $(\sigma_{\psi,z}, \sigma_{\xi,z})$, (π_z^{EU}, π_z^{UE}) , and b_{UI} are chosen to match annual log *household*-income-growth quantiles for comparable lower-income, lower-education households (De Nardi

et al., 2020)—the relevant object for our cohort, rather than the individual earnings that Guvenen et al. (2021) measure directly. We target

$$(P_1, P_5, P_{10}, P_{50}, P_{90}) = (-1.65, -0.79, -0.33, 0.00, 0.20),$$

with annual dispersion $\sigma(\Delta y) = 0.43$. The deep left tail is generated primarily by unemployment spells rather than permanent-income scarring: Guvenen et al. (2021) show that the persistence asymmetry reverses by income level—for low-income workers, negative shocks are relatively transitory—so routing the downside through recoverable unemployment rather than permanent shocks is the empirically appropriate choice, and Section 4 shows it is first-order for the model’s wealth predictions. In the baseline, monthly job-loss probabilities are 0.011 in good states and 0.017 in bad states, monthly job-finding probabilities are 0.11 and 0.09, and the replacement rate is $b_{UI} = 0.15$. The permanent and transitory monthly standard deviations are 0.05–0.06 and 0.10–0.13, respectively.

This calibration uses the bottom-of-distribution evidence in Guvenen et al. (2021) as discipline for the income process of our high-bias cohort rather than as an exact cohort-specific estimate. The maintained empirical premise is that the high-bias households are lower-income, lower-education, and low-saving. Section 4 considers robustness variants that change one element of the income process at a time: loading the downside onto permanent income, fattening the transitory-income tail, or raising b_{UI} . These variants have little effect on the forecast-error swing—shifting unemployment-exit hazards so that P_5 ranges from -0.65 to -0.90 moves the swing only from 0.053 to 0.057—while the wealth level is more sensitive, median liquid wealth moving from 3.7 to 7.0 months.

(3) The asset structure is disciplined by the balance sheet. The liquid returns are set to their empirical values for the cohort: a zero real return on transaction balances, $R_s = 1.00$, and credit-card borrowing at $R_b = 1.18$. The illiquid premium $R_k - R_s$ of 2–4 percent—the illiquid asset is the one that *does* earn a real return—and the withdrawal cost ϕ in the two-asset model are disciplined by the cohort’s liquid/illiquid wealth, *not* by any forecast-error moment.

Consistent with the warning that a large illiquid premium violates stationarity (Section A), the premium is kept moderate and the illiquidity ϕ , rather than the premium, does the work.

(4) The behavioral parameters are calibrated last. Only the behavioral triple $(\beta^E; \beta_G, \beta_B)$ is calibrated to behavioral moments—the two forecast errors FE_G and FE_B , under the ergodic-mean anchor of paragraph (1). The financial-condition pattern $FC_B < FC_G < 0$, the full wealth distribution, debt incidence, and marginal propensities to consume are not targeted and serve as validation. The behavioral targets are $FE_G = 0.044$ and $FE_B = 0.100$ (a swing of 0.056), together with the qualitative pattern $FC_B < FC_G < 0$. The cohort balance-sheet targets are median liquid wealth near zero and median net worth of roughly 8–12 months of income, held largely in illiquid form—the hand-to-mouth balance sheet that low-income, low-education households are documented to hold and that motivates the two-asset structure of Kaplan and Violante (2014). Table 7 summarizes the parameters.

4 Results

The models are solved numerically, and Appendix A gives the details of the solution.

4.1 A constant naivete ratio does not generate the cyclical forecast error

The first and central result concerns the source of the good–bad forecast-error swing. We solve the validated one-asset model and compute FE_G and FE_B from (10). Table 8 reports the case of a state-independent naivete ratio. In the main rows, both sides are constant, with $\beta = 0.7$ and $\beta_G^E = \beta_B^E$, so the ratio β^E/β does not vary with the aggregate state.

The model generates positive average consumption forecast errors: households consume more than they forecast because the direct self-control wedge dominates the endogenous wealth correction from saving less during the year. Thus, the constant-ratio model gets the average sign of the consumption and financial-condition errors right. But it does not generate cyclical-ity.

Holding the realized income process and other primitives fixed, a constant naivete ratio makes the policy wedge the same object in good and bad states; the remaining state dependence comes only through the endogenous wealth correction, which is small. Across values of β^E , the good–bad swing is therefore within Monte Carlo error of zero. A state-independent naivete ratio can match at most one of FE_G and FE_B .

This conclusion is about the ratio, not the levels. Table 8 varies β^E holding β fixed, but the result also holds when both β_z and β_z^E vary across states while their ratio remains constant. Setting $\beta_z^E = \kappa\beta_z$ with β_z equal to the baseline pair, the simulated swing is at most 0.004 in magnitude across $\kappa \in \{1.10, 1.16, 1.25\}$, and has the opposite sign from the data, an order of magnitude below the targeted swing of 0.056. The small residual reflects the endogenous-wealth correction, not a cyclical naivete effect. Conditional on the maintained structure, the forecast-error swing therefore requires state dependence in the naivete wedge.

Table 9 shows that a state-dependent naivete ratio reproduces the joint empirical signature under all three representations. In the baseline representation, true self-control deteriorates from $\beta_G = 0.754$ in good states to $\beta_B = 0.606$ in bad states, while the household continues to expect $\beta^E = 0.81$ of its future selves. The model matches FE_G and FE_B and, without targeting financial-condition forecasts, produces $FC_B < FC_G < 0$. The validation is not the negative financial-condition error by itself: given excess consumption relative to forecast, a negative wealth error follows partly from the budget constraint. The discriminating pattern is the joint sign and cyclicity,

$$FE_B > FE_G > 0, \quad FC_B < FC_G < 0,$$

which the model matches using only the consumption forecast-error moments.

The calibrated values have a simple interpretation. In good states, households are mildly over-optimistic about future self-control, expecting $\beta^E = 0.81$ against realized $\beta_G = 0.754$. In bad states, they fail to anticipate most of the deterioration, expecting the same $\beta^E = 0.81$ against realized $\beta_B = 0.606$. Naivete is therefore countercyclical as an outcome even though beliefs about future self-control are constant: the mistake is not optimism about the external state, but a failure to internalize what bad states do to self-control.

Panel A of Figure 1 summarizes the identification. A state-independent ratio delivers one forecast error per parametrization and therefore lies on the diagonal $FE_G = FE_B$. The empirical target lies off that diagonal, so only state dependence in the naivete ratio can reach it. The same panel also shows the limitation of the survey moments: they identify the state dependence of the ratio, not whether the movement comes from true self-control, beliefs about future self-control, or both. Under the alternative rate-based construction, the spending rate at the twelve-month horizon has partly mean-reverted, and the calibrated baseline delivers a rate swing of only 0.013; the qualitative identification is unchanged. We calibrate to annual spending growth because that is the survey object.

Is the implied deterioration too large? The bad-state value $\beta_B = 0.61$ is a sizeable drop from $\beta_G = 0.75$. Because this state variation has no direct experimental counterpart, we view the magnitude as a calibrated implication rather than as an independently identified estimate. The implied drop is nevertheless consistent with evidence that financial strain impairs cognitive control and shortens planning horizons (Mani et al., 2013; Haushofer and Fehr, 2014). More important, the paper’s central conclusion does not hinge on the exact value: the identifying requirement is state dependence in the naivete ratio, not a particular β_B . Re-solving the model with a higher anchor, $E_z[\beta_z] = 0.75$, raises the bad-state value to $\beta_B = 0.65$ while leaving the forecast-error swing and the financial-condition ordering intact (Section 4.4).

The representations separate in actual behavior, not in forecasts. The three rows of Table 9 are nearly equivalent for the survey moments because forecast errors identify the naivete ratio β_z^E/β_z , not which side of the ratio moves. They differ sharply, however, in actual behavior. In the baseline representation, bad states reduce true self-control and make MPCs countercyclical: the actual-policy monthly MPC at the median household rises from 0.022 in good states to 0.027 in bad states, or roughly 23 versus 28 percent at an annual horizon. About half of this gap comes from the self-control channel and half from countercyclical income risk. By contrast, in the belief-side representation, the same forecast-error moments coexist with essentially acyclical actual behavior. State-dependent beliefs alone are therefore mostly a forecast phenomenon,

whereas state-dependent true self-control changes the behavior that matters for balance-sheet transmission.

Panel B of Figure 1 shows this separation. The three representations stack at the same forecast-error swing, but imply different countercyclical MPC gaps. The belief-side representation sits near the income-risk floor traced by the both-constant benchmark, while the baseline lifts the MPC gap above it. Thus, the survey moments alone do not identify which side of the naivete ratio moves, but the baseline is the representation in which the identified wedge is behaviorally consequential.

4.2 Liquid over-saving is robust in direction but sensitive to income persistence

The second result concerns the wealth implications of the one-asset model. Starting from the quantile-disciplined recoverable-unemployment baseline, we perturb one element of the income process at a time—adding permanent mixture jumps of the size estimated by Guvenen et al. (2021), fattening the transitory tail, replacing the recoverable downside with a Gaussian permanent random walk of similar annual dispersion, and raising the UI floor—and solve the one-asset model under the baseline preference calibration, $(\beta^E, \beta_G, \beta_B) = (0.81, 0.754, 0.606)$. Table 10 reports the resulting income statistics, the growth-impatience modulus of Section A, median liquid wealth, and forecast-error swing.

Three facts emerge. First, the forecast-error swing is stable across the interior income variants: it is 0.055 in the baseline and transitory-tail cases and 0.064 with permanent jumps. It is attenuated, not amplified, in the strong-UI case, falling to 0.037, because the binding constraint compresses the consumption wedge. Thus, the identification result does not depend on the income decomposition, and borrowing constraints work against the swing rather than generate it.

Second, liquid wealth is highly sensitive to income persistence. Permanent income risk moves the model close to, or beyond, the growth-impatience boundary: permanent mixture jumps put the modulus within 4×10^{-4} of the boundary, and a Gaussian permanent random walk of similar

annual dispersion eliminates the stationary normalized wealth distribution altogether. At the other extreme, a generous UI floor removes much of the precautionary motive and collapses households to the borrowing limit. This is not a success for the one-asset model: it lowers liquid wealth by putting essentially everyone in debt, not by producing the liquid-poor but net-worth-positive balance sheet observed in the cohort.

Third, the recoverable-unemployment baseline delivers a moderate liquid buffer of about 5.2 months. Even though it has the deepest annual downside, with $P_{10} = -0.33$, the downside is temporary, so a modest buffer is sufficient to smooth through it. The main lesson is that persistence and insurance, not annual downside depth alone, govern liquid saving.

4.3 What the one-asset model can and cannot match

The one-asset model therefore delivers the main forecast-error results but not the cohort balance sheet. With a state-dependent naivete ratio, it matches the level and cyclicity of consumption forecast errors and produces the financial-condition ordering $FC_B < FC_G < 0$, robustly across income decompositions. But because the only asset is liquid, any precautionary buffer is liquid by construction.

This is the model’s key balance-sheet limitation. Under the recoverable-unemployment baseline, the one-asset model implies about 5.2 months of liquid wealth, while the cohort target is near zero liquid wealth and roughly 8–12 months of net worth held largely in illiquid form. Whether 5.2 months constitutes a sharp failure turns on the cohort’s measured liquid holdings, a dependence we make explicit in Section 5. Changing the income process does not solve this problem: Table 10 shows that the alternatives deliver either an unbounded liquid buffer or debt at the borrowing limit. The missing ingredient is therefore not a different income process but a different asset structure. Section 4.4 introduces a liquid/illiquid asset model to preserve the forecast-error mechanism while matching the cohort’s balance sheet.

4.4 The two-asset model: calibrated results

We now report the calibrated two-asset model under the solver of Section A. Two parameters govern the portfolio block: the illiquid return R_k and the proportional withdrawal cost ϕ . We set $R_k = 1.04$, an illiquid premium at the upper end of the conventional 2–4 percent range that keeps $\delta R_k < 1$, and $\phi = 0.15$. These parameters target the cohort balance sheet: median liquid wealth near zero and median net worth inside the 8–12 month range, held largely in illiquid form. The model delivers a median liquid position of 0.11 months of income, a median illiquid position of 9.4 months, and median net worth of 9.5 months; 46 percent of households hold liquid debt. Thus the one-asset model’s balance-sheet limitation is resolved: the precautionary buffer moves into the return-dominant illiquid asset, while small liquid balances and short borrowing spells absorb month-to-month shocks.

The forecast-error mechanism carries over almost unchanged. Re-estimating the behavioral triple under the same ergodic anchor, $\mathbb{E}_z[\beta_z] = 0.70$, gives

$$(\beta^E; \beta_G, \beta_B) = (0.816; 0.758, 0.607),$$

compared with $(0.81; 0.754, 0.606)$ in the one-asset model. The model matches the targeted consumption forecast errors, $FE_G = 0.044$ and $FE_B = 0.100$, and hence the swing of 0.056. It also preserves the untargeted financial-condition ordering. On liquid wealth, $FC_G = -1.7$ and $FC_B = -4.8$ percent of annual income; on net worth, $FC_G = -4.2$ and $FC_B = -9.6$ percent. Because the survey financial-condition question is broad, the net-worth error is the primary model counterpart. In both cases the disciplined claim is the sign and ordering, $FC_B < FC_G < 0$, not the exact level.

The identification result is also unchanged. Re-solving the two-asset model with a state-independent naivete ratio, $\beta = 0.70$ and $\beta^E = 0.80$, leaves the balance sheet essentially unchanged but produces $FE_G = FE_B = 0.056$: the good–bad swing is zero to four decimal places. As in the one-asset model, asset structure alone does not generate the cyclical forecast-error pattern; state dependence in the naivete ratio is still required.

The second asset adds a behavioral margin the one-asset model cannot express: flows into and out of commitment-like wealth. In the baseline, monthly deposits into the illiquid asset fall by 95 percent in bad states, from 0.044 to 0.002 per unit of monthly permanent income, while withdrawals rise from 0.047 to 0.082. Net flows therefore fall from -0.004 in good states to -0.080 in bad states. Part of this cyclicality is income-driven: under the constant-ratio benchmark, deposits fall by 75 percent and net flows fall from -0.018 to -0.058 . The remaining collapse is behavioral.

To isolate anticipation, we also solve the fully sophisticated benchmark with the same state-dependent true bias and correct beliefs, $\beta_z^E = \beta_z$. This household has zero forecast errors by construction, but differs sharply on the commitment margin. It holds a larger illiquid stock, with median illiquid wealth of 10.6 months, and sustains bad-state deposits of 0.007 per month, more than three times the naive household's 0.002, while good-state deposits are nearly identical. Withdrawal rates per unit of the illiquid stock are similar across the two households, 0.0085 versus 0.0088 per month, so the larger sophisticated outflow in levels reflects a larger stock being drawn down, not a weaker taste for commitment. Thus bad states reduce commitment inflows for all households, but unanticipated deterioration in self-control makes naive households cut those inflows much more sharply.

This is the main new prediction of the two-asset model. Commitment-like saving flows should be especially cyclical for naive households, and the cyclicality of these flows helps distinguish state-dependent self-control from belief-side representations and from projection bias over income, which mainly affect forecasts rather than actual commitment demand. Natural empirical counterparts include retirement-account contributions, mortgage prepayments, and deposits into other illiquid or commitment savings vehicles. Figure 2 displays this margin directly. The constant-ratio benchmark shows the part of the cut attributable to common income risk; the gap between the naive and sophisticated benchmarks isolates the role of anticipation.

Actual-behavior implications are also amplified. Over the ergodic cross-section, the mean actual MPC rises from 0.031 in good states to 0.038 in bad states, and the median doubles from 0.018 to 0.036. Under the constant-ratio benchmark, the corresponding gaps are only 0.001

and 0.002. Near-zero liquid balances therefore magnify the countercyclical MPCs generated by state-dependent true bias relative to the one-asset benchmark.

The results are not a numerical or portfolio-parameter knife edge. Across simulation seeds, additional value iterations, and a refinement of the illiquid grid, the forecast-error swing is unchanged to five decimal places and the near-zero liquid median moves by less than 0.03 months of income. Changing portfolio parameters affects the portfolio split, but not the behavioral block: at $R_k = 1.03$, the liquid and illiquid medians move to (0.92, 6.8) months while forecast errors remain essentially unchanged; at $\phi = 0.10$, the magnitude of the deposit collapse is smaller, but its sign does not change.

Finally, the mechanism is robust to the externally anchored level of present bias. Recalibrating the two-asset model at $\mathbb{E}_z[\beta_z] \in \{0.65, 0.70, 0.75\}$ returns behavioral triples

$$(0.758; 0.704, 0.563), \quad (0.816; 0.758, 0.607), \quad (0.875; 0.812, 0.650),$$

with the same implied naivete ratios to four decimal places. The forecast moments pin the ratios; the anchor sets the level. Median net worth moves from 7.4 to 9.5 to 11.9 months, while the forecast-error swing and the financial-condition ordering remain intact. Figure 3 shows this separation: patience determines the level of wealth, whereas forecast errors identify the mechanism. Table 11 collects the calibration and moments.

5 Discussion: the two comparisons

5.1 One asset versus two assets

The one-asset and two-asset models play distinct roles. The one-asset model identifies the behavioral mechanism. It is the transparent environment in which the central result is cleanest: only a state-dependent naivete ratio generates the cyclical forecast-error swing (Tables 8–9). It also reveals the balance-sheet limitation: once income risk is disciplined, the precautionary buffer appears as liquid wealth (Table 10).

The two-asset model addresses that limitation. We do not claim that a liquid/illiquid asset structure is the only possible fix; differences in credit access or measured consumption needs could also affect liquid saving. But the liquid/illiquid distinction is the standard and minimal way to generate the hand-to-mouth balance sheet emphasized by Kaplan and Violante (2014): little liquid wealth together with positive net worth. The two-asset model is therefore not additional machinery for fitting forecast errors. It preserves the same state-dependent naivete-ratio mechanism while matching the cohort’s balance sheet.

This comparison is conditional on both the data and the income process. The one-asset model’s prediction of about 5.2 months of liquid wealth is a sharp failure only if the relevant cohort is indeed close to liquid-wealth constrained while holding positive illiquid wealth, and the comparison must be made on the disciplined recoverable-unemployment process, since under defensible permanent-risk variants the one-asset liquid prediction is unbounded (Table 10). Under that premise, the two-asset model resolves the balance-sheet problem: at $(R_k, \phi) = (1.04, 0.15)$ it delivers a liquid median of 0.11 months and net worth of 9.5 months, while preserving the forecast-error mechanism, the financial-condition ordering $FC_B < FC_G < 0$, the zero constant-ratio swing, and countercyclical MPCs.

5.2 Constant versus state-dependent naivete

The comparison between constant and state-dependent naivete is the paper’s identifying contrast. Within the recursive partially naive model, a state-independent naivete ratio does not generate the observed state dependence in forecast errors. A state-dependent ratio does. This is the content of Tables 8 and 9.

The identification carries over to the two-asset model. With two assets, the household has an additional deposit/withdrawal margin that could in principle create state asymmetry even with a constant naivete ratio. Quantitatively it does not: re-solving the two-asset model with both β and β^E constant produces $FE_G = FE_B$ to four decimal places (Table 11). The asset structure changes the balance sheet and adds commitment-flow predictions, but it does not replace the need for state dependence in the naivete ratio.

A note on the calibrated behavioral magnitudes. The ergodic mean of β_z is anchored externally, but the state variation in true self-control and the belief level β^E are inferred from the forecast errors. In the baseline, true self-control falls to $\beta_B = 0.61$ in bad states while households continue to expect $\beta^E = 0.81$ of their future selves, compared with a mild good-state gap of 0.81 versus $\beta_G = 0.75$. The identification of the mechanism does not rest on these exact numbers: the qualitative requirement is state dependence in the naivete ratio. The magnitudes should be interpreted as calibrated implications to be assessed against external evidence and, where possible, cross-validated.

6 General equilibrium implications

This section shows that the same state-dependent self-control risk that generates the household forecast errors also has general-equilibrium pricing implications. The exercise is deliberately stripped down: log utility and an endowment economy let us isolate self-control risk from aggregate consumption risk.

We consider a unit mass of identical, infinitesimal agents. The aggregate parameter governing time inconsistency is exogenous and satisfies $\beta_t \in (0, \bar{\beta}]$, where $\bar{\beta} \in (0, 1]$, and the standard exponential discount factor is $\delta \in (0, 1)$. Lower β_t means stronger present bias. Each agent believes that all of her future selves will act with a constant present-bias parameter $\beta^E \in [\bar{\beta}, 1]$, so $\beta^E \geq \beta_t$ in every aggregate state. Agents therefore underestimate the time inconsistency of their own future selves.

The key belief assumption is asymmetric but deliberate. Each agent correctly anticipates the present bias of all other agents and therefore has rational expectations about aggregate consumption, prices, and returns, while misperceiving only her own future behavior. The continuum economy makes this diagnostic separation possible: an infinitesimal agent's mistake about herself has no effect on aggregate equilibrium quantities.

Let $W_{i,t}$ denote agent i 's beginning-of-period wealth, including the current dividend. The

self who acts at date t chooses current consumption and portfolio weights $\omega_{i,t}$ and solves

$$\max_{C_{i,t}, \omega_{i,t}} \left\{ \log C_{i,t} + \beta_t \mathbb{E}_t \sum_{j=1}^{\infty} \delta^j \log \widehat{C}_{i,t+j} \right\}, \quad (15)$$

subject to

$$W_{i,t+1} = (W_{i,t} - C_{i,t})R_{i,t}. \quad (16)$$

Here $R_{i,t}$ is the gross return from date t to date $t+1$ generated by $\omega_{i,t}$, and $\widehat{C}_{i,t+j}$ is the consumption that the agent expects her own future self to choose. Future own choices are forecast as if future selves used β^E , while aggregate choices are forecast using the true future realizations β_{t+j} .

Aggregate consumption C_t is an exogenous endowment. For simplicity, C_{t+1}/C_t is i.i.d. and independent of β_t .

With log utility, the equilibrium wealth-consumption ratio is

$$\phi(\beta_t) \equiv \frac{W_t}{C_t} = 1 + \frac{\beta_t \delta}{1 - \delta}, \quad (17)$$

while the wealth-consumption ratio that an agent assigns to her own future self is

$$\phi^E = 1 + \frac{\beta^E \delta}{1 - \delta}. \quad (18)$$

Appendix B gives the derivations.

The stochastic discount factor is

$$\begin{aligned} M_{t+1} &= \frac{\phi(\beta_t) - 1}{\phi(\beta_{t+1}) - 1} \frac{C_t}{C_{t+1}} \\ &= \frac{\beta_t \delta}{1 - \delta + \beta_{t+1} \delta} \frac{C_t}{C_{t+1}}, \end{aligned} \quad (19)$$

and the return on the aggregate consumption claim is

$$\begin{aligned} R_{C,t+1} &= \frac{\phi(\beta_{t+1})}{\phi(\beta_t) - 1} \frac{C_{t+1}}{C_t} \\ &= \frac{1 - \delta + \beta_{t+1}\delta}{\beta_t\delta} \frac{C_{t+1}}{C_t}. \end{aligned} \quad (20)$$

Thus prices and returns vary with the future realization of β_{t+1} even though agents have rational expectations about aggregate quantities. The perceived parameter β^E governs forecasts of own future consumption, but it does not enter aggregate prices in this log-utility benchmark.

To isolate the pricing of self-control risk, set aggregate consumption risk to zero, so $C_{t+1}/C_t = 1$. Then

$$M_{t+1} = \frac{\phi(\beta_t) - 1}{\phi(\beta_{t+1})}, \quad R_{C,t+1} = \frac{\phi(\beta_{t+1})}{\phi(\beta_t) - 1}. \quad (21)$$

The conditional risk premium on the consumption claim is

$$\begin{aligned} \mathbb{E}_t[R_{C,t+1}] - R_{f,t} &= \frac{1}{\phi(\beta_t) - 1} \left\{ \mathbb{E}_t[\phi(\beta_{t+1})] - \frac{1}{\mathbb{E}_t[1/\phi(\beta_{t+1})]} \right\} \\ &= \frac{1}{\beta_t\delta} \left\{ 1 - \delta + \delta\mathbb{E}_t[\beta_{t+1}] - \frac{1}{\mathbb{E}_t[1/(1 - \delta + \beta_{t+1}\delta)]} \right\} > 0, \end{aligned} \quad (22)$$

whenever β_{t+1} is conditionally non-degenerate. The inequality follows because the arithmetic mean of $\phi(\beta_{t+1})$ exceeds its harmonic mean. The consumption claim therefore carries a positive risk premium even though aggregate consumption is risk-free.

The intuition is direct. A low realization of β_{t+1} makes other agents consume a larger fraction of wealth, lowering the payoff on the aggregate consumption claim. The agent correctly prices this aggregate payoff risk. At the same time, her own planned future consumption is based on the perceived rule indexed by β^E , so the low wealth delivered in low- β_{t+1} states is a low perceived-consumption, high-marginal-utility state. The consumption claim pays poorly precisely when perceived marginal utility is high. The premium is time-varying because the current valuation ratio and, if β_t is persistent, the conditional distribution of β_{t+1} vary over time.

The same distinction between rational aggregate expectations and biased expectations about own behavior generates predictable individual consumption forecast errors. Since all agents hold

the aggregate consumption claim in equilibrium, actual individual consumption growth equals aggregate consumption growth:

$$\frac{C_{i,t+1}}{C_{i,t}} = \frac{C_{t+1}}{C_t}. \quad (23)$$

The agent correctly forecasts next-period wealth, but believes that her future self will consume the fraction $1/\phi^E$ of that wealth. Her forecast of own consumption growth is therefore

$$F_{i,t} \left(\frac{C_{i,t+1}}{C_{i,t}} \right) = \mathbb{E}_t \left[\frac{\phi(\beta_{t+1})}{\phi^E} \frac{C_{t+1}}{C_t} \right]. \quad (24)$$

The expected consumption-growth forecast error is

$$\begin{aligned} \mathbb{E}_t \left[\frac{C_{i,t+1}}{C_{i,t}} \right] - F_{i,t} \left(\frac{C_{i,t+1}}{C_{i,t}} \right) &= \mathbb{E} \left[\frac{C_{t+1}}{C_t} \right] \left\{ 1 - \frac{\mathbb{E}_t[\phi(\beta_{t+1})]}{\phi^E} \right\} \\ &= \mathbb{E} \left[\frac{C_{t+1}}{C_t} \right] \frac{\delta (\beta^E - \mathbb{E}_t[\beta_{t+1}])}{1 - \delta + \beta^E \delta}. \end{aligned} \quad (25)$$

The error is positive whenever $\beta^E > \mathbb{E}_t[\beta_{t+1}]$. If β_t is persistent, a low current realization predicts a low future realization, so agents underpredict their own consumption growth most strongly in states in which actual self-control is expected to remain weak. They do so while continuing to forecast aggregate consumption and asset prices rationally.

The same mechanism causes agents to overestimate their future financial conditions. To match the empirical convention in the paper, define the financial-condition forecast error as realized financial conditions minus expected financial conditions. The date- t self correctly forecasts beginning-of-period wealth $W_{i,t+1}$, but expects the next self to consume $W_{i,t+1}/\phi^E$ rather than the larger amount $W_{i,t+1}/\phi(\beta_{t+1})$. Hence

$$\begin{aligned} &\mathbb{E}_t [W_{i,t+1} - C_{i,t+1}] - F_{i,t} (W_{i,t+1} - C_{i,t+1}) \\ &= \mathbb{E}_t \left[W_{i,t+1} \left(\frac{1}{\phi^E} - \frac{1}{\phi(\beta_{t+1})} \right) \right] \\ &= (1 - \delta) \mathbb{E}_t \left[W_{i,t+1} \left(\frac{1}{1 - \delta + \beta^E \delta} - \frac{1}{1 - \delta + \beta_{t+1} \delta} \right) \right] \leq 0. \end{aligned} \quad (26)$$

The inequality is strict whenever $\beta_{t+1} < \beta^E$ with positive conditional probability. Thus, excessive

optimism about future financial conditions appears as a negative forecast error: realized wealth after consumption falls short of expected retained wealth. Because forecast and actual selves choose the same portfolio weights under log utility, the discrepancy is carried forward by the same realized returns. Repeated overconsumption relative to plan therefore compounds the overprediction of wealth at longer horizons; Appendix B.6 gives the multi-period expression.

6.1 Rational benchmark with time-varying discounting

A rational time-varying discount-factor model does not generate the same implications. Consider the same representative-agent endowment economy, but replace quasi-geometric discounting with time-consistent preferences and a Markov exponential discount factor δ_t :

$$U_{i,t}^R = \log C_{i,t} + \delta_t \mathbb{E}_t [U_{i,t+1}^R] . \quad (27)$$

Agents know the process for δ_t and have rational expectations about their own future preferences, consumption, and all aggregate quantities.

Let $\phi^R(\delta_t)$ denote the wealth-consumption ratio. With log utility,

$$\phi^R(\delta_t) = 1 + \delta_t \mathbb{E}_t [\phi^R(\delta_{t+1})] . \quad (28)$$

The stochastic discount factor and the return on the consumption claim are

$$M_{t+1}^R = \delta_t \frac{C_t}{C_{t+1}} , \quad (29)$$

and

$$R_{C,t+1}^R = \frac{\phi^R(\delta_{t+1})}{\phi^R(\delta_t) - 1} \frac{C_{t+1}}{C_t} . \quad (30)$$

When aggregate consumption is constant, $M_{t+1}^R = \delta_t$ is known at date t , so

$$R_{f,t}^R = \frac{1}{\delta_t} . \quad (31)$$

Using (28),

$$\begin{aligned}\mathbb{E}_t [R_{C,t+1}^R] &= \frac{\mathbb{E}_t [\phi^R(\delta_{t+1})]}{\phi^R(\delta_t) - 1} \\ &= \frac{1}{\delta_t} = R_{f,t}^R.\end{aligned}\tag{32}$$

Thus,

$$\mathbb{E}_t [R_{C,t+1}^R] - R_{f,t}^R = 0.\tag{33}$$

The return on the consumption claim can vary with δ_{t+1} , but this variation is not priced because the stochastic discount factor is conditionally deterministic when aggregate consumption is constant.

The rational benchmark also generates no predictable individual consumption forecast errors:

$$F_{i,t}^R \left(\frac{C_{i,t+1}}{C_{i,t}} \right) = \mathbb{E}_t \left[\frac{C_{i,t+1}}{C_{i,t}} \right],\tag{34}$$

and therefore

$$\mathbb{E}_t \left[\frac{C_{i,t+1}}{C_{i,t}} \right] - F_{i,t}^R \left(\frac{C_{i,t+1}}{C_{i,t}} \right) = 0.\tag{35}$$

With constant aggregate consumption, both actual and expected individual consumption growth equal one. Time-varying rational impatience can move valuations, but it does not generate self-control risk premia or predictable own-consumption forecast errors in this benchmark.

7 Model implications and extensions

The endowment economy isolates the implications of time-varying present bias for asset prices and household forecast errors. The same force may also affect environments in which households face costly adjustment decisions.

7.1 State-dependent monetary-policy transmission

One implication concerns monetary-policy transmission. Maxted et al. (2025) show that naive present bias creates a motive to procrastinate on mortgage refinancing, an immediate-cost, delayed-benefit action. This delay reduces the impact response of consumption to an interest-rate cut and prolongs pass-through to household spending.

In their Calvo-style adjustment mechanism, let $\bar{\varepsilon}$ denote the usual high effort cost and let $\varepsilon < \bar{\varepsilon}$ denote the cost during a temporary low-cost opportunity. A present-biased household completes the adjustment during that opportunity when

$$\varepsilon < \beta_t \bar{\varepsilon}. \quad (36)$$

Defining $\beta^* \equiv \varepsilon/\bar{\varepsilon}$, suppose

$$\beta_B < \beta^* < \beta_G. \quad (37)$$

A household with $\beta_t = \beta_G$ refinances when the low-cost opportunity arrives, whereas a household with $\beta_t = \beta_B$ continues to procrastinate. If low realizations of β_t occur in bad aggregate states, the refinancing channel of monetary policy operates more slowly precisely in downturns. This statement concerns the timing of transmission. Stronger present bias may still raise marginal propensities to consume and amplify the eventual spending response once refinancing occurs.

A simple extension avoids the implication that bad-state households postpone refinancing forever. Let $\bar{\varepsilon}$ be the normal cost, ε_M a moderately low cost, and ε_L a very low cost, with

$$\varepsilon_L < \beta_B \bar{\varepsilon} < \varepsilon_M < \beta_G \bar{\varepsilon}. \quad (38)$$

Both households procrastinate at the normal cost. At ε_M , only the household with $\beta_t = \beta_G$ refinances; at ε_L , both households refinance. If the medium- and very-low-cost opportunities arrive at rates λ_M and λ_L , the adjustment hazards conditional on being in the refinancing region are

$$\lambda_G = \lambda_M + \lambda_L, \quad \lambda_B = \lambda_L. \quad (39)$$

Time-varying present bias therefore generates a state-dependent propagation lag in monetary-policy transmission. This mechanism is outside the current endowment economy, but it illustrates how the same state-dependent self-control risk can matter in household balance-sheet adjustment.

8 Conclusion

Households do not only forecast aggregate conditions; they also forecast their own future behavior. Using the New York Fed Survey of Consumer Expectations, we show that these personal forecasts contain systematic, opposite-signed errors: realized consumption growth exceeds expected consumption growth, while realized personal financial conditions fall short of expectations. Both errors are concentrated among households plausibly more exposed to self-control problems and widen during economic stress. In bad times, these households underpredict consumption growth even more while becoming more optimistic about their future finances. The joint pattern points to expected restraint that does not materialize: households plan to consume cautiously and end up financially better off, but subsequently consume more and improve their finances less than anticipated.

We interpret these facts through a model of state-dependent partial naiveté. Present-biased agents underestimate how much their future selves will overconsume, and this mistake becomes more severe in bad states. The mechanism is not misforecasting aggregate conditions. Agents may forecast the external environment correctly; the error is a misforecast of their own future behavior, which leaves them over-optimistic about their future finances.

The model accounts for the joint consumption and financial-condition forecast errors. A constant naivete ratio can match at most the average sign of the errors; it does not generate the good–bad swing. A state-dependent naivete ratio does. Adding liquid and illiquid assets preserves the mechanism, matches the cohort’s balance sheet, and generates predictions for commitment-like saving flows, such as retirement-account contributions, mortgage prepayments, and other illiquid saving.

The broader message is that predictable errors in personal consumption expectations reveal state-varying self-control problems. Because these errors are systematic, they can matter beyond

survey responses: they affect saving, balance sheets, and, in equilibrium, discount rates. Future work could use richer balance-sheet and transaction data to test the mechanism more directly.

References

- Andries, M., T. M. Eisenbach, and M. C. Schmalz. 2024. Horizon-Dependent Risk Aversion and the Timing and Pricing of Uncertainty. *Review of Financial Studies* 37, 3272–3334.
- Arnsten, A. F. T. 2009. Stress Signaling Pathways that Impair Prefrontal Cortex Structure and Function. *Nature Reviews Neuroscience* 10, 410–422.
- Baumeister, C., D. Leiva-León, and E. Sims. 2024. Tracking Weekly State-Level Economic Conditions. *Review of Economics and Statistics* 106, 483–504.
- Carroll, C. D. 1997. Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis. *Quarterly Journal of Economics* 112, 1–55.
- De Nardi, M., G. Fella, and G. Paz-Pardo. 2020. Nonlinear Household Earnings Dynamics, Self-Insurance, and Welfare. *Journal of the European Economic Association* 18, 890–926.
- DellaVigna, S., and U. Malmendier. 2004. Contract Design and Self-Control: Theory and Evidence. *Quarterly Journal of Economics* 119, 353–402.
- Eliasz, K., and R. Spiegel. 2006. Contracting with Diversely Naive Agents. *Review of Economic Studies* 73, 689–714.
- Fedyk, A. 2025. Asymmetric Naiveté: Beliefs about Self-Control. *Management Science* 71, 6047–6068.
- Federal Reserve Bank of New York. 2024. Survey of Consumer Expectations, Household Spending Survey. *Median Year-Ahead Growth in Overall Household Spending*. <https://www.newyorkfed.org/microeconomics/sce>.
- Gabaix, X., and D. Laibson. 2022. Myopia and Discounting. NBER working paper.
- Groneck, M., A. Ludwig, and A. Zimmer. 2024. Who Saves More, the Naive or the Sophisticated Agent? *Journal of Economic Theory* 219, 1–33.

- Guvenen, F., F. Karahan, S. Ozkan, and J. Song. 2021. What Do Data on Millions of U.S. Workers Reveal about Lifecycle Earnings Dynamics? *Econometrica* 89, 2303–2339.
- Harris, C., and D. Laibson. 2001. Dynamic Choices of Hyperbolic Consumers. *Econometrica* 69, 935–957.
- Harris, C., and D. Laibson. 2013. Instantaneous Gratification. *Quarterly Journal of Economics* 128, 205–248.
- Haushofer, J., and E. Fehr. 2014. On the Psychology of Poverty. *Science* 344, 862–867.
- Heidhues, P., and B. Köszegi. 2010. Exploiting Naivete about Self-Control in the Credit Market. *American Economic Review* 100, 2279–2303.
- Kaplan, G., and G. L. Violante. 2014. A Model of the Consumption Response to Fiscal Stimulus Payments. *Econometrica* 82, 1199–1239.
- Khapko, M. 2023. Asset Pricing with Dynamically Inconsistent Agents. *Finance and Stochastics* 27, 1017–1046.
- Krusell, P., and A. A. Smith. 2003. Consumption-Savings Decisions with Quasi-Geometric Discounting. *Econometrica* 71, 365–375.
- Laibson, D. 1997. Golden Eggs and Hyperbolic Discounting. *Quarterly Journal of Economics* 112, 443–477.
- Luttmer, E. G. J., and T. Mariotti. 2003. Subjective Discounting in an Exchange Economy. *Journal of Political Economy* 111, 959–989.
- Mani, A., S. Mullainathan, E. Shafir, and J. Zhao. 2013. Poverty Impedes Cognitive Function. *Science* 341, 976–980.
- Maxted, P. 2025. Present Bias Unconstrained: Consumption, Welfare, and the Present-Bias Dilemma. *Quarterly Journal of Economics* 140, 2963–3013.

- Maxted, P., D. Laibson, and B. Moll. 2025. Present Bias Amplifies the Household Balance-Sheet Channels of Macroeconomic Policy. *Quarterly Journal of Economics* 140, 691–743.
- O'Donoghue, T., and M. Rabin. 1999. Doing It Now or Later. *American Economic Review* 89, 103–124.
- O'Donoghue, T., and M. Rabin. 2001. Choice and Procrastination. *Quarterly Journal of Economics* 116, 121–160.
- Phelps, E. S., and R. A. Pollak. 1968. On Second-Best National Saving and Game-Equilibrium Growth. *Review of Economic Studies* 35, 185–199.
- Sapolsky, R. 2017. *Behave: The Biology of Humans at Our Best and Worst*. London: Penguin Press.
- Strotz, R. H. 1956. Myopia and Inconsistency in Dynamic Utility Maximization. *Review of Economic Studies* 23, 165–180.

Table 1: Summary statistics

The table reports the summary statistics of the New York Fed’s Survey of Consumer Expectations raw, unfiltered data. Detailed explanations of the variables are provided in Section 2.1.

	Mean	St. dev.	N	Min	P1	P99	Max
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Δc	1.12	63.59	27,819	-477.47	-41.94	42.83	9,997.86
$F[\Delta c]$	-1.12	90.11	156,756	-11,190	-65	50	24,965
Education level	0.28	0.34	157,041	0	0	1	1
Income level	0.53	0.40	155,920	0	0	1	1
Age	50.54	15.54	157,399	0	23	82	891
Risk tolerance	0.39	0.30	121,145	0	0	1	1
Not save generally	0.66	0.47	24,604	0	0	1	1

Table 2: Consumption growth forecast errors and demographics

The table reports the difference between individuals' realized and expected consumption growth based on the New York Fed's Survey of Consumer Expectations data. Column (1) presents the average consumption growth forecast error based on the full sample period (December 2014 until December 2022), while Column (2) focuses on the Pre-Covid (pre-2020) period. Columns (3) and (4) report the estimates of a regression with time and state fixed effects of consumption growth forecast errors on demographic characteristics: education, income, age, risk tolerance (each classified in three normalized categories 0, 0.5 and 1, consistent with the New York Fed's categorization), and an indicator of not saving in general. Individuals with the lowest level of education and income are assigned the value of 1. The realized and expected consumption growth are trimmed at 1% by time. Individuals between 25 and 80 years old are included in the final sample of 11,928 unique individuals. The t-statistics, reported in brackets below, are based on Driscoll-Kraay standard errors using 5 lags. Significance at 10%, 5% and 1% is denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)
Education			3.461*** (3.725)	3.367*** (3.516)
Income			1.727*** (4.591)	1.732*** (5.288)
Not save generally			0.357*** (4.156)	0.341*** (2.911)
Age				0.293 (0.972)
Risk tolerance				-0.529* (-1.904)
Constant	1.675*** (5.980)	2.059*** (7.257)		
Time and State FE	N	N	Y	Y
Driscoll-Kraay SE	Y	Y	Y	Y
Pre-Covid period	N	Y	N	N
R-squared	0.000	0.000	0.019	0.019
N	25,817	16,494	23,398	22,529

Table 3: Consumption growth and financial situation forecast errors

The table reports estimates based on individuals' forecast errors from the New York Fed's Survey of Consumer Expectations data for the sample period from December 2014 until December 2022. Columns (1) and (2) use consumption growth forecast errors, defined as the difference between realized and expected consumption growth; Columns (3) and (4) use financial situation forecast errors. Panel A reports unconditional forecast errors: Columns (1) and (3) report the average forecast error (constant only, with Driscoll–Kraay standard errors using 5 lags), while Columns (2) and (4) report the estimates of a regression with time and state fixed effects of forecast errors on the bias cohort, a composite of education, income, and an indicator of not saving in general, estimated in the Online Appendix. Panel B reports the estimates of a regression with time and state fixed effects of average forecast errors on the bias cohort, a stress indicator, and the interaction between cohort and stress. In Columns (1) and (3) the stress variable $z_{j,t-1}$ is a state-level economic conditions indicator (as in Baumeister et al. (2024)), and in Columns (2) and (4) it is the state-level change in unemployment compared to the year before (available in FRED). The t-statistics in Panel B are based on standard errors clustered by cohort and state. The realized and expected consumption growth are trimmed at 1% by time, and individuals between 25 and 80 years old are included in the final sample of 11,928 unique individuals. The t-statistics are reported in brackets below the estimates. Significance at 10%, 5% and 1% is denoted by *, **, and ***, respectively.

	Consumption growth FE		Financial situation FE	
	(1)	(2)	(3)	(4)
<i>Panel A: Unconditional forecast errors</i>				
Bias cohort		5.545*** (4.446)		-0.202*** (-7.417)
Constant	1.675*** (5.980)		-0.129*** (-9.832)	
Time and State FE	N	Y	N	Y
Driscoll–Kraay SE	Y	Y	Y	Y
R-squared	0.000	0.019	0.000	0.016
N	25,817	23,398	25,799	23,385
<i>Panel B: Time variation in forecast errors</i>				
Bias cohort	6.283*** (10.795)	6.250*** (10.809)	-0.170*** (-4.104)	-0.166*** (-3.988)
Stress indicator	-0.422 (-1.360)	-0.348 (-1.504)	0.016 (0.651)	0.013 (0.595)
Cohort × Stress indicator	0.896** (2.482)	0.775** (2.134)	-0.054** (-2.297)	-0.040* (-1.771)
Time and State FE	Y	Y	Y	Y
Clustered SE	Y	Y	Y	Y
R-Squared	0.049	0.048	0.071	0.070
N	5,549	5,557	5,547	5,555

Table 4: Consumption expectations and financial distress

The table reports the relationship between the bias cohort and household consumption expectations and financial distress outcomes using data from the New York Fed’s Survey of Consumer Expectations from December 2014 to December 2022. The bias cohort is constructed from education, income, and saving behavior as described in the Online Appendix. Panel A reports estimates for expected non-essential consumption growth, expected essential consumption growth, and the difference between the two. Panel B reports estimates for indicators of maxing out credit cards, being late on loan payments by more than 30 days, and being late on loan payments by more than 90 days during the previous year. All regressions include time and state fixed effects. The reported t-statistics are based on Driscoll–Kraay standard errors with five lags. Significance at the 10%, 5%, and 1% levels is denoted by *, **, and ***, respectively.

<i>Panel A: Consumption expectations</i>			
	Non-essential (1)	Essential (2)	Noness-Essential (3)
Bias cohort	-10.427*** (-9.119)	-3.690*** (-8.853)	-6.110*** (-8.944)
Time and State FE	Y	Y	Y
Driscoll–Kraay SE	Y	Y	Y
R-Squared	0.045	0.018	0.033
N	23,114	23,161	22,923
<i>Panel B: Financial Distress</i>			
	Maxed-out credit card (1)	Late loan pmt (30 Days) (2)	Late loan pmt (90 Days) (3)
Bias cohort	0.157*** (11.246)	0.098*** (8.429)	0.310*** (5.905)
Time and State FE	Y	Y	Y
Driscoll–Kraay SE	Y	Y	Y
R-Squared	0.026	0.023	0.103
N	12,397	15,071	836

Table 5: Time variation in non-essential and essential consumption expectations

The table reports the estimates of a regression with time and state fixed effects of non-essential and essential consumption growth expectations on bias cohorts (composed from education, income and not saving in general indicators and estimated the online appendix), a stress indicator, and the interaction between cohort and stress. In Panel A the stress variable $z_{j,t-1}$ is a state-level change in unemployment compared to the year before (available in FRED) and in Panel B – a state-level economic conditions indicator (as in Baumeister et al. (2024)). In Column (1) and (2) the dependent variable is non-essential and essential consumption growth expectations over the next 12 months, respectively. In Column (3) the dependent variable is the difference between non-essential and essential consumption growth expectations. The expected and realized consumption growth along with demographic characteristics are available in the New York Fed’s Survey of Consumer Expectations data for the sample period from December 2014 until December 2022. The realized and expected consumption growth are trimmed at 1% by time and individuals between 25 and 80 years old are included in the final sample of 11,928 unique individuals. The t-statistics, reported in brackets below, are based on clustered by cohort and state standard errors. Significance at 10%, 5% and 1% is denoted by *, **, and ***, respectively.

	Non-essential (1)	Essential (2)	Noness-Essential (3)
<i>Panel A: State-level Unemployment</i>			
Cohort	-10.437*** (-15.743)	-4.400*** (-10.308)	-5.267*** (-11.117)
Stress indicator	0.451 (1.596)	0.015 (0.109)	0.363 (1.505)
Cohort × Stress indicator	-1.421*** (-3.516)	-0.554*** (-2.603)	-0.608** (-2.008)
Time and State FE	Y	Y	Y
Clustered SE	Y	Y	Y
R-Squared	0.096	0.048	0.059
N	5,527	5,531	5,508
<i>Panel B: State-level Economic Index</i>			
Cohort	-10.515*** (-15.814)	-4.423*** (-10.347)	-5.316*** (-11.174)
Stress indicator	0.229 (0.684)	0.115 (0.579)	0.162 (0.589)
Cohort × Stress indicator	-1.723*** (-4.206)	-0.584*** (-2.742)	-0.830*** (-2.728)
Time and State FE	Y	Y	Y
Clustered SE	Y	Y	Y
R-Squared	0.099	0.048	0.060
N	5,520	5,524	5,502

Table 6: Financially unconstrained sample: Time variation in consumption growth forecast errors

The table focuses on a subsample of financially unconstrained consumers. Detailed definition of the financially constrained consumers indicator is provided in Section 2.5. The table reports the estimates of a regression with time and state fixed effects of average consumption growth forecast errors based on bias buckets, a stress indicator, and the interaction between bias and stress. In Columns (1) and (2) the stress indicator is a state-level economic conditions indicator (as in Baumeister et al. (2024)) and in Columns (3) and (4) – a state-level change in unemployment compared to the year before (available in FRED). In Columns (1) and (3) the bias is estimated based on education, income and an indicator of not saving in general, while in Columns (2) and (4) the bias is estimated based on education and the indicator of not saving in general. The expected and realized consumption growth along with demographic characteristics used to estimate the bias buckets are available in the New York Fed’s Survey of Consumer Expectations data for the sample period from December 2014 until December 2022. The realized and expected consumption growth are trimmed at 1% by time and individuals between 25 and 80 years old are included in the final sample of 7,546 unique financially unconstrained individuals. The t-statistics, reported in brackets below, are based on clustered by cohort and state standard errors. Significance at 10%, 5% and 1% is denoted by *, **, and ***, respectively.

	Change in unemployment		Economic conditions	
	Bias 1 (1)	Bias 2 (2)	Bias 1 (3)	Bias 2 (4)
Bias	6.206*** (8.070)	4.705*** (6.320)	6.294*** (8.124)	4.687*** (6.373)
Stress indicator	-0.188 (-0.565)	-0.240 (-0.725)	0.099 (0.239)	-0.300 (-0.681)
Bias $\times$ Stress indicator	1.428*** (2.980)	0.956** (2.382)	1.306*** (2.769)	1.149*** (2.976)
Time and State FE	Y	Y	Y	Y
Clustered SE	Y	Y	Y	Y
R-Squared	0.061	0.063	0.060	0.063
N	2,990	2,327	2,987	2,319

Table 7: Parameters

Symbol	Description	Value
β_G, β_B	true present bias by state; <i>anchored</i> $\mathbb{E}_z[\beta_z] = 0.70$	0.754, 0.606
δ	long-run discount (annual; monthly $\delta_m = \delta^{1/12}$)	0.96
γ	relative risk aversion	2
R_s, R_b	liquid save / borrow gross real return (annual; monthly $R^{1/12}$)	1.00, 1.18
R_k, ϕ	illiquid return (annual) / withdrawal cost (two-asset)	$R_s + 2-4\%$, 0.15
Π	aggregate-state chain; $P(B)$ stationary	$[[.969, .031], [.050, .950]]$; 0.38
income	permanent RW + transitory + employment; recoverable downside	see §3.4
β^E	perceived future bias, constant; <i>calibrated</i> (with β_z) to FE_G, FE_B	0.81

Notes: δ, γ , the income quantiles, and the ergodic-mean anchor $\mathbb{E}_z[\beta_z] = 0.70$ are fixed before any behavioral moment is examined. Given the anchor, $(\beta^E; \beta_G, \beta_B)$ are just-identified by the two forecast-error targets. The illiquid block (R_k, ϕ) is disciplined by the cohort balance sheet.

Table 8: A constant naivete ratio generates no cyclical swing

β^E (constant)	FE_G	FE_B	swing $FE_B - FE_G$	FC_G, FC_B
0.76	0.035	0.034	-0.000	-3.2%, -3.1%
0.80	0.057	0.056	-0.001	-5.1%, -4.9%
0.84	0.077	0.076	-0.001	-6.9%, -6.6%
0.88	0.097	0.096	-0.001	-8.6%, -8.3%

Notes: One-asset model, $\beta = 0.7$, $\gamma = 2$; forecast error is the percentage year-ahead spending-growth object (10); FC is the financial-condition error as a fraction of annual income. The swing is within twice the Monte Carlo standard error of zero at every value, and the financial-condition error is acyclical. The result is identical under the alternative rate-based construction. Income process: the quantile-disciplined recoverable-unemployment baseline of Section 3.7. The median household holds about five months of liquid wealth, so a non-negligible share of paths interact with the borrowing constraint; the swing remains within Monte Carlo error of zero.

Table 9: Three representations of a state-dependent naive ratio match the joint signature

Representation	FE_G	FE_B	swing	FC_G, FC_B	liq. (mo.)	ΔMPC
(1) beliefs vary	0.044	0.099	0.055	-4.0, -8.5	5.4	+0.002
(2) baseline (β_z varies)	0.044	0.099	0.055	-3.9, -8.8	5.2	+0.005
(3) constant share	0.044	0.099	0.055	-4.0, -9.2	4.6	+0.010

Notes: One-asset model, $\gamma = 2$, annual-spending-growth FE; FC entries are percent of annual income. Targets are $FE_G = 0.044$, $FE_B = 0.100$ (swing 0.056); the financial-condition pattern is untargeted in every row. *Parameters:* (1) $\beta = 0.70$, $\beta_z^E = (0.755, 0.930)$; (2) $\beta^E = 0.81$, $\beta_z = (0.754, 0.606)$; (3) $\lambda = 0.66$, $\beta_z = (0.830, 0.491)$, so $\beta_z^E = (0.888, 0.666)$. Representations (2)–(3) impose the ergodic-mean anchor $\mathbb{E}_z[\beta_z] = 0.70$; (1) fixes $\beta = 0.70$. Hitting the targets requires $\beta_B = 0.49$ in (3), outside the experimental range, and β_B^E near one in (1), infeasible for $\gamma \geq 3$. ΔMPC is the actual-policy monthly MPC, bad minus good state, at the median household ($m = 5.5$ months); the both-constant benchmark (countercyclical income risk only) delivers +0.002. Under the alternative rate-based construction the swing attenuates sharply (0.013 in the baseline; 0.017 in representation (1), where no $\beta_B^E \leq 1$ can reach the empirical magnitude); the calibration matches the annual-growth construction, which is the survey object.

Table 10: One-asset liquid wealth across income decompositions

Income process	$\sigma(\Delta y)$	P_{10}	GIC mod.	liquid (mo.)	FE swing
Recoverable unemployment (baseline)	0.43	−0.33	0.9986	5.2	0.055
+ permanent mixture jumps	0.44	−0.36	0.9996	n.s.	0.064
+ transitory fat tail	0.43	−0.33	0.9986	5.1	0.055
Gaussian permanent random walk	0.45	−0.37	1.0010	none	—
Strong UI floor ($b_{UI} = 0.45$)	0.20	−0.19	0.9986	−1.6 [†]	0.037

Notes: One-asset model under the calibrated baseline ($\beta^E = 0.81$; $\beta_G, \beta_B = 0.754, 0.606$); each row perturbs one dimension of the income process. The GIC modulus is the bad-state growth-impatience factor $(R_s^{1/12} \delta_m)^{1/\gamma} \mathbb{E}[\psi^{-1}]$ of Section A; values at or above one imply no stationary normalized wealth distribution. “n.s.”: the simulated liquid distribution fails to settle at any burn-in length (even its 10th percentile rises without bound), reflecting a modulus within 4×10^{-4} of the boundary. “none”: the modulus exceeds one, so no stationary distribution exists and no swing is reported. †: in the strong-UI row the precautionary motive collapses and every household is in debt (median −1.6 months against the limit $\underline{a} = -3$); the bad-state swing is also *attenuated* there (0.037), because the binding constraint compresses the consumption wedge—constraints attenuate the swing, they do not create it. The forecast-error swing is stable across the interior rows. The cohort target for median liquid wealth is approximately zero.

Table 11: The two-asset model: calibration, moments, and behavioral benchmarks

	Baseline	Constant ratio	Sophisticated
<i>Portfolio block (calibrated to the balance sheet)</i>			
Illiquid return R_k (annual)	1.04	1.04	1.04
Withdrawal cost ϕ	0.15	0.15	0.15
<i>Behavioral block (anchored at $\mathbb{E}_z[\beta_z] = 0.70$)</i>			
Beliefs (β_G, β_B)	$\beta^E = 0.816$ (0.758, 0.607)	$\beta^E = 0.80$ $\beta = 0.70$	$\beta_z^E = \beta_z$ (0.758, 0.607)
<i>Targeted moments</i>			
FE_G (target 0.044)	0.044	0.056	0
FE_B (target 0.100)	0.100	0.056	0
Swing (target 0.056)	0.056	0.000	0
Liquid median, months (target ≈ 0)	0.11	0.23	0.06
Net worth median, months (target 8–12)	9.5	9.7	10.7
<i>Untargeted moments and diagnostics</i>			
Illiquid median, months	9.4	9.2	10.6
Debt incidence	0.46	0.42	0.47
FC liquid (G, B), pct. of annual income	(−1.7, −4.8)	(−1.8, −2.1)	0
FC net worth (G, B)	(−4.2, −9.6)	(−5.4, −5.2)	0
Deposits per month (G, B)	(0.044, 0.002)	(0.035, 0.009)	(0.042, 0.007)
Withdrawals per month (G, B)	(0.047, 0.082)	(0.052, 0.066)	(0.050, 0.090)
Mean MPC (G, B), ergodic	(0.031, 0.038)	(0.030, 0.031)	(0.033, 0.062)
Median MPC (G, B), ergodic	(0.018, 0.036)	(0.020, 0.022)	(0.022, 0.041)
ΔMPC (B–G), reference point	+0.012	+0.001	—

Notes: production grid $N_m = 130$, $N_k = 44$; perceived-value residual of order 10^{-5} for the baseline and constant-ratio columns and 10^{-4} for the sophisticated column (whose flow moments are stable under further iteration; see text). Moments are stable across simulation seeds, iteration depth, and an illiquid grid of $N_k = 56$ to within one percent (the swing to five decimal places); the liquid median is stable to within 0.03 months of income, and the bad-state deposit flow is 0.002 on both grids. The constant-ratio and sophisticated columns re-solve the model at identical (R_k, ϕ) ; the sophisticated column sets Markov-consistent correct beliefs $\beta_z^E = \beta_z$ at the baseline (β_G, β_B) , so its forecast and financial-condition errors are zero by construction and its content is the commitment margin. Deposit and withdrawal flows are monthly means per unit of monthly permanent income under the actual policy over the ergodic distribution. MPCs are central differences over one quarter of a month of income; the reference point is one month of liquid resources at the median illiquid position.




Figure 1: Identification and its corollary (one asset, $\gamma = 2$). *Panel A:* in the (FE_G, FE_B) plane every state-independent naive ratio lies on the diagonal $FE_G = FE_B$ (connected markers sweep a constant β^E at $\beta = 0.70$); the empirical target lies off the diagonal, so a state-dependent ratio is required to reach it, and all three representations of Table 9 coincide on that point. *Panel B:* each representation plotted by its forecast-error swing $FE_B - FE_G$ against the countercyclical MPC gap ΔMPC_{B-G} at the median household. The representations are indistinguishable on the swing but separate in actual behavior; the belief-side representation lies on the income-risk floor set by the both-constant benchmark, while the baseline (state-dependent true bias) lifts the gap above it.




Figure 2: Commitment inflows collapse in bad states (two asset). Monthly deposits into the illiquid asset by aggregate state, under the baseline (naive, state-dependent true bias), the constant ratio, and the fully sophisticated benchmark (correct beliefs, same true bias). All three cut inflows in bad states, but the naive household—failing to anticipate the deterioration—cuts most sharply (−95%), while the sophisticated household facing the identical true bias sustains roughly three times the bad-state inflow. Withdrawal rates per unit of stock are similar across the naive and sophisticated columns (text), so the diagnostic contrast is on the inflow margin.




Figure 3: Level versus identification (two asset). As the externally anchored ergodic mean of true bias $\mathbb{E}_z[\beta_z]$ rises from 0.65 to 0.75, median net worth (left axis) moves through the empirical 8–12 month band, while the bad-state forecast error and the cyclical swing (right axis) are essentially unchanged and the implied naivete ratios are identical across anchors to four decimal places. Patience sets the level; the forecast moments identify the mechanism.

A Numerical method

One-asset model: time iteration with the present-bias envelope. We solve the perceived (β_z^E) problem by time iteration on the endogenous gridpoint method. The object iterated is the perceived marginal continuation value implied by (6): by the envelope condition for quasi-hyperbolic consumers (Harris and Laibson, 2001, see also Maxted, 2025),

$$\widehat{V}_m = u'(\widehat{c}) \left[\frac{1}{\beta_z^E} - \text{MPC} \left(\frac{1}{\beta_z^E} - 1 \right) \right], \quad (\text{A.1})$$

where MPC is the marginal propensity to consume out of cash-on-hand along the forecast policy (the formula collapses to the standard envelope $\widehat{V}_m = u'(\widehat{c})$ where the borrowing constraint binds, since there $\text{MPC} = 1$). The actual policy (7) is then a single one-shot-deviation step, with the true bias β_z in force at the decision date, against the converged $\widehat{V}$.

Equilibrium selection and the growth-impatience boundary. Two well-known features of discrete-time quasi-hyperbolic models shape the method. First, sophisticated quasi-geometric consumers admit a continuum of discontinuous, step-function Markov equilibria (Krusell and Smith, 2003); the economically standard object is the *smooth* equilibrium, the limit of finite-horizon backward induction. Policy-improvement algorithms can converge to the discontinuous equilibria, so we compute the smooth equilibrium by iterating the backward-induction map itself, accelerated with Anderson mixing, and verify smoothness of the converged policy directly (no downward step in consumption exceeds 3×10^{-3} on the solution grid). Second, at monthly frequency the normalized model sits close to the growth-impatience boundary: the contraction modulus of the time iteration is $(R_s^{1/12} \delta_m)^{1/\gamma} \mathbb{E}[\psi^{-1}]$, which equals 1.0002 at a conventional 4 percent liquid real return—there the normalized time-consistent benchmark has no finite buffer-stock target and the iteration does not converge—and 0.9986 at our calibrated $R_s = 1.00$. The zero liquid real return is the empirically appropriate value for this cohort, and it is also the value at which the model is well posed. Perceived present bias adds effective impatience through the hyperbolic Euler discount $\delta_m[\beta_z^E + (1 - \beta_z^E) \text{MPC}']$ (Harris and Laibson, 2001), strengthening the contraction further. The same boundary reappears in the extreme income-risk variants of

Table 10, where we report the modulus alongside each specification.

Validation. The solver is validated as follows: (i) at $\beta = \beta^E$ the actual policy coincides with the sophisticated policy to 4×10^{-11} , and the converged policies are smooth and monotone; (ii) at high wealth, where the endogenous-response correction vanishes, the ratio $c^*/\hat{c}$ converges to the closed-form wedge (8): at $\gamma = 2$ and constant $\beta^E = 0.9$ it is 1.1335 at $m = 300$ months against $(0.9/0.7)^{1/2} = 1.1339$, approaching the benchmark monotonically in wealth and lying below it at moderate wealth, exactly as the endogenous-response caveat of Section 3 predicts; (iii) median normalized Euler-equation residuals at off-grid points are 1.7×10^{-5} (forecast policy) and 1.2×10^{-5} (actual), with maxima of 5×10^{-3} adjacent to the borrowing kink; (iv) the simulated state-conditional forecast errors are tightly bracketed by the occupancy-weighted static wedges implied by the aggregate chain (0.052 and 0.120 predicted against 0.044 and 0.099 simulated), with the simulated values *below* the frictionless benchmark and the gap larger in the bad state, where the borrowing constraint binds more often—the sign and ranking the endogenous-response terms imply; and (v) the headline moments are stable to grid refinement (the swing moves by 10^{-4} from 240 to 480 gridpoints and across an extended grid with twice the wealth range), across simulation seeds (spread 6×10^{-4}), and across burn-in lengths (1,500 versus 3,000 months, 2×10^{-4}). The rebaselined specification (state-dependent β_z against a constant β^E) inherits all of these properties; its seed spread is 7×10^{-4} . Three caveats are stated for completeness: the extended-grid solve terminates at a sup-norm residual of order 10^{-2} (its moments are nonetheless within 10^{-4} of baseline); the near-boundary permanent-jump variant of Table 10 terminates at 5×10^{-4} (its modulus is within 4×10^{-4} of one and its wealth distribution is non-stationary by construction); and the $\gamma = 1$ benchmark wedge approaches its closed form ($0.9/0.7 \approx 1.286$) monotonically but has not fully attained it at the largest wealth we report. All headline results—the $\gamma = 2$ baseline grid—meet the strict 10^{-10} criterion.

The impatience condition in the two-asset model. The same boundary disciplines the illiquid return. Because present bias discounts only the immediate period, long-run accumulation of the illiquid asset is governed by the annual product δR_k (invariant to compounding frequency):

if $\delta R_k > 1$, the time-consistent ($\beta = 1$) household accumulates the illiquid asset without bound and no stationary distribution exists. Present bias lowers the effective long-run discount factor toward $\delta R_k [\beta + (1 - \beta) \text{MPC}']$ (Harris and Laibson, 2001) and can restore stationarity, but relying on that would be a warning rather than a virtue: it would indicate that R_k is too high relative to δ . With the liquid return at its empirical value $R_s = 1.00$, an illiquid premium of 2–4 percent keeps $\delta R_k < 1$, so the two-asset model is stationary under conventional discounting and the illiquidity ϕ carries the second asset’s role—exactly as the calibration strategy requires.

For the two-asset model we use value-function iteration on the perceived continuation value. Within each candidate illiquid target, the liquid choice is solved by an endogenous-grid step; the deposit, no-trade, and withdrawal regimes are then compared by direct value maximization across the full set of illiquid targets. Because no first-order condition is inverted across regimes, the comparison is robust to the kink non-concavity that the withdrawal cost introduces at the no-trade boundary—the upper-envelope failures of nested endogenous-grid methods, which made the liquid moment unstable in earlier passes, do not arise. Anderson acceleration is applied to the value map. The perceived-value residual converges to the order of 10^{-5} , and we report the two-asset balance sheet only after verifying the stated stability criterion: across simulation seeds, additional value iterations, and a refinement of the illiquid grid, the liquid and illiquid moments must be stable to within one percent. Section 4.4 reports the battery.

B Derivations for the log-utility economy

B.1 Consumption rules

To derive the individual policies, let $\widehat{V}_t(W)$ denote the continuation value generated by the sequence of future selves that the agent believes will use β^E . With log utility, conjecture

$$\widehat{V}_t(W) = a \log W + b_t, \tag{B.2}$$

where b_t can depend on the aggregate state through equilibrium returns. The forecast policy of a future self solves

$$\max_C \left\{ \log C + \beta^E \delta \mathbb{E}_t \left[\widehat{V}_{t+1}((W - C)R_{i,t}) \right] \right\}. \tag{B.3}$$

The first-order condition is

$$\frac{1}{C} = \frac{\beta^E \delta a}{W - C}, \quad (\text{B.4})$$

which implies

$$\frac{W}{\widehat{C}} = 1 + \beta^E \delta a. \quad (\text{B.5})$$

The continuation value is evaluated using standard exponential discounting between future dates,

$$\widehat{V}_t(W) = \log \widehat{C} + \delta \mathbb{E}_t \left[\widehat{V}_{t+1}((W - \widehat{C})R_{i,t}) \right]. \quad (\text{B.6})$$

Substitution of (B.2) into (B.6) and matching the coefficient on $\log W$ gives

$$a = 1 + \delta a = \frac{1}{1 - \delta}. \quad (\text{B.7})$$

Equation (B.5) therefore becomes

$$\frac{W}{\widehat{C}} = 1 + \frac{\beta^E \delta}{1 - \delta} = \phi^E. \quad (\text{B.8})$$

The self who actually acts at date t uses the true current time-inconsistency parameter β_t while evaluating the future with the same perceived continuation value. Her consumption choice solves

$$\max_C \left\{ \log C + \beta_t \delta \mathbb{E}_t \left[\widehat{V}_{t+1}((W - C)R_{i,t}) \right] \right\}. \quad (\text{B.9})$$

The corresponding first-order condition gives

$$\frac{W}{C} = 1 + \beta_t \delta a = 1 + \frac{\beta_t \delta}{1 - \delta} = \phi(\beta_t). \quad (\text{B.10})$$

The forecast and actual portfolio first-order conditions differ only by a positive multiplicative constant. Hence, with log utility, the forecast and actual selves choose the same portfolio conditional on prices; partial naivete affects the consumption-saving margin but not the portfolio composition directly.

B.2 Equilibrium prices and the pricing kernel

In a symmetric equilibrium, market clearing implies

$$W_t = \phi(\beta_t)C_t. \quad (\text{B.11})$$

Let P_t denote the ex-dividend price of the aggregate consumption claim. Since $W_t = P_t + C_t$,

$$\frac{P_t}{C_t} = \phi(\beta_t) - 1 = \frac{\beta_t \delta}{1 - \delta}. \quad (\text{B.12})$$

The gross return on the consumption claim is

$$\begin{aligned} R_{C,t+1} &= \frac{P_{t+1} + C_{t+1}}{P_t} = \frac{W_{t+1}}{W_t - C_t} \\ &= \frac{\phi(\beta_{t+1})}{\phi(\beta_t) - 1} \frac{C_{t+1}}{C_t} \\ &= \frac{1 - \delta + \beta_{t+1}\delta}{\beta_t \delta} \frac{C_{t+1}}{C_t}, \end{aligned} \quad (\text{B.13})$$

which gives (20).

For a one-period state-contingent payoff, the first-order condition from (B.9) gives

$$M_{t+1} = \beta_t \delta \frac{\widehat{V}'_{t+1}(W_{i,t+1})}{1/C_{i,t}} = \frac{\beta_t \delta}{1 - \delta} \frac{C_{i,t}}{W_{i,t+1}}. \quad (\text{B.14})$$

In the symmetric equilibrium,

$$\frac{W_{i,t+1}}{C_{i,t}} = \phi(\beta_{t+1}) \frac{C_{t+1}}{C_t}. \quad (\text{B.15})$$

Using $\beta_t \delta / (1 - \delta) = \phi(\beta_t) - 1$ in (B.14) yields

$$\begin{aligned} M_{t+1} &= \frac{\phi(\beta_t) - 1}{\phi(\beta_{t+1})} \frac{C_t}{C_{t+1}} \\ &= \frac{\beta_t \delta}{1 - \delta + \beta_{t+1}\delta} \frac{C_t}{C_{t+1}}, \end{aligned} \quad (\text{B.16})$$

where the second equality uses

$$\phi(\beta_t) - 1 = \frac{\beta_t \delta}{1 - \delta}, \quad \phi(\beta_{t+1}) = \frac{1 - \delta + \beta_{t+1}\delta}{1 - \delta}. \quad (\text{B.17})$$

This is (19).

B.3 Risk premium without aggregate consumption risk

When $C_{t+1}/C_t = 1$, the risk-free rate satisfies

$$R_{f,t}^{-1} = \mathbb{E}_t[M_{t+1}] = (\phi(\beta_t) - 1) \mathbb{E}_t \left[\frac{1}{\phi(\beta_{t+1})} \right]. \quad (\text{B.18})$$

The conditional expected return on the consumption claim is

$$\mathbb{E}_t[R_{C,t+1}] = \frac{\mathbb{E}_t[\phi(\beta_{t+1})]}{\phi(\beta_t) - 1}. \quad (\text{B.19})$$

Subtracting the risk-free rate gives (22). Since, for any positive non-degenerate random variable X ,

$$\mathbb{E}_t[X] > \frac{1}{\mathbb{E}_t[1/X]}, \quad (\text{B.20})$$

the premium is strictly positive whenever $\phi(\beta_{t+1})$, and hence β_{t+1} , is conditionally non-degenerate.

B.4 Rational benchmark with time-varying discounting

For existence and boundedness, it is sufficient to assume that there is a constant $\bar{\delta} < 1$ such that $\delta_t \leq \bar{\delta}$ almost surely. This condition ensures that the expected discounted utility sum is finite and that the recursion in (28) has a unique bounded solution. The finite-state assumption is not required.

For the rational benchmark, conjecture that the value function takes the form

$$V_t^R(W) = a(\delta_t) \log W + b(\delta_t). \quad (\text{B.21})$$

The consumption first-order condition implies

$$\frac{W}{C} = 1 + \delta_t \mathbb{E}_t[a(\delta_{t+1})]. \quad (\text{B.22})$$

Matching the coefficient on $\log W$ in the Bellman equation gives

$$a(\delta_t) = 1 + \delta_t \mathbb{E}_t[a(\delta_{t+1})]. \quad (\text{B.23})$$

Thus, $a(\delta_t) = \phi^R(\delta_t)$ and the wealth-consumption ratio satisfies (28). Under the boundedness

condition above, iterating the recursion forward gives

$$\phi^R(\delta_t) = 1 + \mathbb{E}_t \left[\sum_{j=1}^{\infty} \prod_{s=0}^{j-1} \delta_{t+s} \right], \quad (\text{B.24})$$

which is finite because the product is bounded above by $\bar{\delta}^j$.

The Euler equation for any one-period return R_{t+1} is

$$1 = \mathbb{E}_t \left[\delta_t \frac{C_t}{C_{t+1}} R_{t+1} \right], \quad (\text{B.25})$$

which yields the stochastic discount factor in (29). Market clearing gives the consumption-claim return in (30). When $C_{t+1}/C_t = 1$, equation (28) implies

$$\phi^R(\delta_t) - 1 = \delta_t \mathbb{E}_t [\phi^R(\delta_{t+1})]. \quad (\text{B.26})$$

Therefore,

$$\mathbb{E}_t [R_{C,t+1}^R] = \frac{\mathbb{E}_t [\phi^R(\delta_{t+1})]}{\delta_t \mathbb{E}_t [\phi^R(\delta_{t+1})]} = \frac{1}{\delta_t} = R_{f,t}^R. \quad (\text{B.27})$$

Finally, because agents know the law of motion for δ_t and correctly anticipate their own future policy, their subjective conditional expectation of consumption growth equals the objective conditional expectation, giving (35).

B.5 Individual consumption forecasts

From (16), (B.10), and (20), individual wealth next period is

$$\begin{aligned} W_{i,t+1} &= (W_{i,t} - C_{i,t}) R_{C,t+1} \\ &= (\phi(\beta_t) - 1) C_{i,t} \frac{\phi(\beta_{t+1})}{\phi(\beta_t) - 1} \frac{C_{t+1}}{C_t} \\ &= C_{i,t} \phi(\beta_{t+1}) \frac{C_{t+1}}{C_t}. \end{aligned} \quad (\text{B.28})$$

The agent forecasts this wealth correctly. She believes, however, that her future self will consume $W_{i,t+1}/\phi^E$. Thus,

$$\frac{\widehat{C}_{i,t+1}}{C_{i,t}} = \frac{\phi(\beta_{t+1})}{\phi^E} \frac{C_{t+1}}{C_t}. \quad (\text{B.29})$$

In reality, the future self consumes $W_{i,t+1}/\phi(\beta_{t+1})$, so

$$\frac{C_{i,t+1}}{C_{i,t}} = \frac{C_{t+1}}{C_t}. \quad (\text{B.30})$$

Taking conditional expectations, using the independence of C_{t+1}/C_t and β_{t+1} , and substituting (17)–(18) gives

$$\begin{aligned} \mathbb{E}_t \left[\frac{C_{i,t+1}}{C_{i,t}} \right] - F_{i,t} \left(\frac{C_{i,t+1}}{C_{i,t}} \right) &= \mathbb{E} \left[\frac{C_{t+1}}{C_t} \right] \left(1 - \frac{\mathbb{E}_t[\phi(\beta_{t+1})]}{\phi^E} \right) \\ &= \mathbb{E} \left[\frac{C_{t+1}}{C_t} \right] \frac{\delta(\beta^E - \mathbb{E}_t[\beta_{t+1}])}{1 - \delta + \beta^E \delta}, \end{aligned} \quad (\text{B.31})$$

which is (25).

B.6 Individual wealth forecasts

The agent correctly forecasts $W_{i,t+1}$ but expects every subsequent self to save the fraction $1 - 1/\phi^E$ of beginning-of-period wealth. Along a given realization of aggregate states and returns, her date- t forecast of wealth h periods after $t + 1$ is therefore

$$\widehat{W}_{i,t+h+1} = W_{i,t+1} \prod_{j=1}^h \left(1 - \frac{1}{\phi^E} \right) R_{i,t+j}. \quad (\text{B.32})$$

Actual wealth evolves according to

$$W_{i,t+h+1} = W_{i,t+1} \prod_{j=1}^h \left(1 - \frac{1}{\phi(\beta_{t+j})} \right) R_{i,t+j}. \quad (\text{B.33})$$

The forecast and actual selves choose the same portfolio weights, so the return $R_{i,t+j}$ is common to both expressions. Since $\beta^E \geq \beta_{t+j}$ implies

$$1 - \frac{1}{\phi^E} \geq 1 - \frac{1}{\phi(\beta_{t+j})}, \quad (\text{B.34})$$

forecast wealth is weakly greater than actual wealth along every realized path. More explicitly,

$$\frac{\widehat{W}_{i,t+h+1}}{W_{i,t+h+1}} = \prod_{j=1}^h \frac{1 - 1/\phi^E}{1 - 1/\phi(\beta_{t+j})} \geq 1. \quad (\text{B.35})$$

Each additional period in which $\beta_{t+j} < \beta^E$ adds another factor greater than one to (B.35). Hence, repeated overconsumption relative to plan compounds the agent's overprediction of future wealth. Under the realized-minus-expected convention used in the empirical analysis, the corresponding h -period financial-condition forecast error satisfies

$$W_{i,t+h+1} - \widehat{W}_{i,t+h+1} \leq 0, \quad (\text{B.36})$$

with a strict inequality along any path containing at least one date at which $\beta_{t+j} < \beta^E$.

C Extension to CRRA utility

This appendix derives the economy under CRRA utility,

$$u(C) = \frac{C^{1-\gamma}}{1-\gamma}, \quad \gamma > 0, \quad \gamma \neq 1. \quad (\text{C.37})$$

We maintain the assumptions that aggregate consumption growth is i.i.d. and independent of the process for β_t . Define

$$\mu_\gamma \equiv \mathbb{E} \left[\left(\frac{C_{t+1}}{C_t} \right)^{1-\gamma} \right], \quad (\text{C.38})$$

and assume that this moment is finite. The log-utility economy is obtained as the limit $\gamma \rightarrow 1$.

C.1 Perceived continuation value and consumption rules

Let $\widehat{V}_t(W)$ denote the continuation value generated by the policies that the agent believes her future selves will follow. Homotheticity implies

$$\widehat{V}_t(W) = A(\beta_t) \frac{W^{1-\gamma}}{1-\gamma}, \quad (\text{C.39})$$

where $A(\beta_t) > 0$. Let $\phi^E(\beta_t)$ denote the wealth-consumption ratio of a future self that the agent believes will use β^E , and let $\phi(\beta_t)$ denote the wealth-consumption ratio of the self that actually acts with β_t .

For a given portfolio opportunity set, define

$$\mathcal{Q}(\beta_t) \equiv \mathbb{E}_t [A(\beta_{t+1}) R_{i,t}^{1-\gamma}] \quad (\text{C.40})$$

evaluated at the optimal portfolio. The portfolio maximizes

$$\frac{1}{1-\gamma} \mathbb{E}_t [A(\beta_{t+1}) R_{i,t}^{1-\gamma}]. \quad (\text{C.41})$$

Because β_t and β^E enter the continuation term only as positive multiplicative constants, the actual and perceived future selves choose the same portfolio conditional on prices.

The perceived future self solves

$$\max_C \left\{ \frac{C^{1-\gamma}}{1-\gamma} + \beta^E \delta \mathcal{Q}(\beta_t) \frac{(W-C)^{1-\gamma}}{1-\gamma} \right\}. \quad (\text{C.42})$$

The first-order condition implies

$$[\phi^E(\beta_t) - 1]^\gamma = \beta^E \delta \mathcal{Q}(\beta_t). \quad (\text{C.43})$$

The self who actually acts at date t solves the same problem with β_t replacing β^E , so

$$[\phi(\beta_t) - 1]^\gamma = \beta_t \delta \mathcal{Q}(\beta_t). \quad (\text{C.44})$$

Combining (C.43) and (C.44) gives

$$\phi^E(\beta_t) - 1 = \left(\frac{\beta^E}{\beta_t} \right)^{1/\gamma} [\phi(\beta_t) - 1]. \quad (\text{C.45})$$

Since $\beta^E \geq \beta_t$, the perceived future self consumes a weakly smaller fraction of wealth:

$$\phi^E(\beta_t) \geq \phi(\beta_t). \quad (\text{C.46})$$

The continuation value satisfies

$$\widehat{V}_t(W) = u\left(\frac{W}{\phi^E(\beta_t)}\right) + \delta \mathbb{E}_t \left[\widehat{V}_{t+1} \left(W \left(1 - \frac{1}{\phi^E(\beta_t)} \right) R_{i,t} \right) \right]. \quad (\text{C.47})$$

Substituting (C.39) and using (C.43) yields

$$A(\beta_t) = [\phi^E(\beta_t)]^{\gamma-1} \left[1 + \frac{\phi^E(\beta_t) - 1}{\beta^E} \right]. \quad (\text{C.48})$$

Equations (C.45) and (C.48) reduce the individual problem to the determination of the actual wealth-consumption ratio $\phi(\beta_t)$.

C.2 General equilibrium

In a symmetric equilibrium, the aggregate consumption claim is the portfolio held by every agent and

$$R_{C,t+1} = \frac{\phi(\beta_{t+1})}{\phi(\beta_t) - 1} \frac{C_{t+1}}{C_t}. \quad (\text{C.49})$$

Using (C.49) in (C.40), together with the independence of aggregate consumption growth and β_{t+1} , gives

$$\mathcal{Q}(\beta_t) = [\phi(\beta_t) - 1]^{\gamma-1} \mu_\gamma \mathbb{E}_t [A(\beta_{t+1}) \phi(\beta_{t+1})^{1-\gamma}]. \quad (\text{C.50})$$

Substitution into (C.44) gives the equilibrium recursion

$$\phi(\beta_t) - 1 = \beta_t \delta \mu_\gamma \mathbb{E}_t [A(\beta_{t+1}) \phi(\beta_{t+1})^{1-\gamma}]. \quad (\text{C.51})$$

The CRRA equilibrium is therefore characterized by the three equations

$$\phi^E(\beta_t) - 1 = \left(\frac{\beta^E}{\beta_t} \right)^{1/\gamma} [\phi(\beta_t) - 1], \quad (\text{C.52})$$

$$A(\beta_t) = [\phi^E(\beta_t)]^{\gamma-1} \left[1 + \frac{\phi^E(\beta_t) - 1}{\beta^E} \right], \quad (\text{C.53})$$

$$\phi(\beta_t) - 1 = \beta_t \delta \mu_\gamma \mathbb{E}_t [A(\beta_{t+1}) \phi(\beta_{t+1})^{1-\gamma}]. \quad (\text{C.54})$$

Unlike the log-utility case, the wealth-consumption ratios generally do not admit a state-by-state closed form. Equations (C.52)–(C.54) form a nonlinear fixed-point system. We assume that the process for β_t and the aggregate-consumption moment in (C.38) are such that this system has a finite positive solution.

C.3 Pricing kernel and the consumption claim

The first-order condition for a one-period state-contingent payoff gives

$$\begin{aligned} M_{t+1} &= \beta_t \delta \frac{\widehat{V}'_{t+1}(W_{i,t+1})}{u'(C_{i,t})} \\ &= \beta_t \delta A(\beta_{t+1}) \left(\frac{C_{i,t}}{W_{i,t+1}} \right)^\gamma. \end{aligned} \quad (\text{C.55})$$

In symmetric equilibrium,

$$M_{t+1} = \beta_t \delta \frac{A(\beta_{t+1})}{\phi(\beta_{t+1})^\gamma} \left(\frac{C_t}{C_{t+1}} \right)^\gamma. \quad (\text{C.56})$$

Using (C.48), this can also be written entirely in terms of the actual and perceived wealth-consumption ratios:

$$M_{t+1} = \beta_t \delta \frac{[\phi^E(\beta_{t+1})]^{\gamma-1} [1 + (\phi^E(\beta_{t+1}) - 1) / \beta^E]}{\phi(\beta_{t+1})^\gamma} \left(\frac{C_t}{C_{t+1}} \right)^\gamma. \quad (\text{C.57})$$

The return on the consumption claim remains given by (C.49).

When aggregate consumption is constant, the stochastic discount factor and the consumption-claim return reduce to

$$M_{t+1} = \beta_t \delta \frac{A(\beta_{t+1})}{\phi(\beta_{t+1})^\gamma}, \quad R_{C,t+1} = \frac{\phi(\beta_{t+1})}{\phi(\beta_t) - 1}. \quad (\text{C.58})$$

Thus, time-varying present bias continues to generate variation in both the pricing kernel and the return even in the absence of aggregate consumption risk. The conditional premium can be written as

$$\mathbb{E}_t[R_{C,t+1}] - R_{f,t} = -\frac{\text{Cov}_t(M_{t+1}, R_{C,t+1})}{\mathbb{E}_t[M_{t+1}]}. \quad (\text{C.59})$$

Equivalently, the premium is positive if

$$\text{Cov}_t \left(\frac{A(\beta_{t+1})}{\phi(\beta_{t+1})^\gamma}, \phi(\beta_{t+1}) \right) < 0. \quad (\text{C.60})$$

For log utility, $A(\beta_{t+1}) = 1/(1 - \delta)$ is constant and the first argument in (C.60) is proportional to $1/\phi(\beta_{t+1})$. The covariance is then strictly negative whenever β_{t+1} is conditionally non-degenerate, recovering the positive-premium result in the main text. Under general CRRA, the sign of the premium depends on the equilibrium response of the continuation-value coefficient as well as the wealth-consumption ratio.

C.4 Consumption and wealth forecast errors

Because all agents hold the aggregate consumption claim in the symmetric equilibrium, actual individual consumption growth equals aggregate consumption growth:

$$\frac{C_{i,t+1}}{C_{i,t}} = \frac{C_{t+1}}{C_t}. \quad (\text{C.61})$$

The agent correctly forecasts next-period wealth but believes that her future self will consume the fraction $1/\phi^E(\beta_{t+1})$ rather than $1/\phi(\beta_{t+1})$. Hence,

$$F_{i,t} \left(\frac{C_{i,t+1}}{C_{i,t}} \right) = \mathbb{E}_t \left[\frac{\phi(\beta_{t+1})}{\phi^E(\beta_{t+1})} \frac{C_{t+1}}{C_t} \right]. \quad (\text{C.62})$$

Using independence,

$$\begin{aligned} & \mathbb{E}_t \left[\frac{C_{i,t+1}}{C_{i,t}} \right] - F_{i,t} \left(\frac{C_{i,t+1}}{C_{i,t}} \right) \\ &= \mathbb{E} \left[\frac{C_{t+1}}{C_t} \right] \left\{ 1 - \mathbb{E}_t \left[\frac{\phi(\beta_{t+1})}{\phi^E(\beta_{t+1})} \right] \right\} \geq 0. \end{aligned} \quad (\text{C.63})$$

The inequality follows from (C.46), and it is strict whenever $\beta_{t+1} < \beta^E$ with positive conditional probability.

Using the same realized-minus-expected convention as in the empirical analysis, the corresponding one-period financial-condition forecast error is

$$\begin{aligned} & \mathbb{E}_t [W_{i,t+1} - C_{i,t+1}] - F_{i,t} (W_{i,t+1} - C_{i,t+1}) \\ &= \mathbb{E}_t \left[W_{i,t+1} \left(\frac{1}{\phi^E(\beta_{t+1})} - \frac{1}{\phi(\beta_{t+1})} \right) \right] \leq 0. \end{aligned} \quad (\text{C.64})$$

Thus, as under log utility, the agent underpredicts her own future consumption and overpredicts the wealth that remains after consumption. Excessive optimism about future financial conditions therefore appears as a negative forecast error.

The discrepancy compounds at longer horizons. Along any realized path of aggregate states

and returns,

$$\widehat{W}_{i,t+h+1} = W_{i,t+1} \prod_{j=1}^h \left(1 - \frac{1}{\phi^E(\beta_{t+j})} \right) R_{i,t+j}, \quad (\text{C.65})$$

$$W_{i,t+h+1} = W_{i,t+1} \prod_{j=1}^h \left(1 - \frac{1}{\phi(\beta_{t+j})} \right) R_{i,t+j}. \quad (\text{C.66})$$

Because the actual and perceived future selves choose the same portfolio and $\phi^E(\beta_{t+j}) \geq \phi(\beta_{t+j})$,

$$\frac{\widehat{W}_{i,t+h+1}}{W_{i,t+h+1}} = \prod_{j=1}^h \frac{1 - 1/\phi^E(\beta_{t+j})}{1 - 1/\phi(\beta_{t+j})} \geq 1. \quad (\text{C.67})$$

Every future period in which the agent overestimates self-control adds another factor greater than one to the forecast-to-actual wealth ratio. Accordingly, the realized-minus-expected financial-condition forecast error satisfies

$$W_{i,t+h+1} - \widehat{W}_{i,t+h+1} \leq 0, \quad (\text{C.68})$$

with strict inequality whenever the perceived and actual saving policies differ along the realized path.

C.5 Rational benchmark under CRRA utility

For completeness, consider the rational benchmark with time-varying exponential discounting and CRRA utility. Conjecture

$$V_t^R(W) = A^R(\delta_t) \frac{W^{1-\gamma}}{1-\gamma}, \quad (\text{C.69})$$

and let $\phi^R(\delta_t)$ denote the wealth-consumption ratio. The consumption first-order condition and the Bellman equation imply

$$A^R(\delta_t) = [\phi^R(\delta_t)]^\gamma. \quad (\text{C.70})$$

Using market clearing and the consumption-claim return gives

$$\phi^R(\delta_t) - 1 = \delta_t \mu_\gamma \mathbb{E}_t [\phi^R(\delta_{t+1})]. \quad (\text{C.71})$$

A sufficient boundedness condition is that $\delta_t \leq \bar{\delta} < 1$ almost surely and

$$\bar{\delta}\mu_\gamma < 1. \quad (\text{C.72})$$

The stochastic discount factor is the standard CRRA kernel

$$M_{t+1}^R = \delta_t \left(\frac{C_t}{C_{t+1}} \right)^\gamma. \quad (\text{C.73})$$

When aggregate consumption is constant, $M_{t+1}^R = \delta_t$ is known at date t . Therefore,

$$R_{f,t}^R = \frac{1}{\delta_t}, \quad \mathbb{E}_t[R_{C,t+1}^R] - R_{f,t}^R = 0. \quad (\text{C.74})$$

Rational expectations also imply zero predictable consumption and wealth forecast errors.