Choice under Fundamental Uncertainty: The Case of Aggregate Consumption

Johannes Binswanger* Albert Flak† Manuel Oechslin‡

May 16, 2022

Abstract

We consider an agent who does not know the data-generating process (DGP) of the economic environment but observes past outcomes. Moreover, the agent considers it possible that the DGP changes in unpredictable ways. In this setting—which we refer to as one of fundamental uncertainty—standard optimal intertemporal choice is not feasible. We provide a model in which the agent makes forward-looking decisions using a future value function that does not depend on any specific information about a DGP. The agent makes forecasts about a subsequent period based on historical analogies. Specifically, we consider the consumption and asset holding decision of a representative agent who earns an exogenous stream of labor income. We calibrate the model to aggregate US data. Despite its simplicity, the model captures the relevant empirical patterns better than a rational expectations model with a comparable degree of flexibility.

JEL classification: D81, D84, E21

Keywords: Fundamental uncertainty, value function, aggregate consumption, expectation formation, historical analogies

*University of St. Gallen, Department of Economics, Switzerland. Email: johannes.binswanger@unisg.ch.

†University of St. Gallen, Department of Economics, Switzerland. Email: albert.flak@gmail.com.

‡University of Lucerne, Department of Economics, Switzerland. Email: manuel.oechslin@unilu.ch.

1 Introduction

It is a common assumption in economic models that agents know the underlying data-generating process (DGP). In fact, this is a defining feature of what we understand as the *rational expectations* paradigm. In dynamic macroeconomic models, this means that agents know the statistical distributions of the fundamental shocks, and how they feed into economic decisions. The rational expectations modeling strategy has been a striking success. One of its major achievements in macroeconomics has been to demonstrate that policy makers cannot systematically manipulate agents' behavior by fooling their expectations. Another major contribution has been to clarify how microeconomic decisions feed into macroeconomic outcomes.

In spite of these successes, the rational expectations assumption also gives raise to some puzzling observations. Maybe most strikingly, while model agents perfectly know the DGP of their model economies, the best informed real-world decision makers—such as central bank rate setting committees and their staff—arguably struggle with identifying the appropriate DGP from incomplete data. This is reflected in countless newspapers articles which feature "sense making" attempts about the DGP behind current developments. As a recent example, consider a quote from *The Economist* (2022):

> Economists are struggling to forecast how many people who left the workforce in 2020 will eventually return [...] They are also grappling with doubts over when consumers will shift their spending back to services, easing the upward pressure on goods prices caused by bunged-up supply chains. Economic data have become harder to interpret. If retail sales fall, for example, does it reflect economic weakening, or a welcome return to normal patterns of consumption?

In our view, the expressions of "struggling to forecast", "grappling with doubts" or "hard to interpret" do not just reflect difficulties of disentangling signal and noise in an observation from a stochastic process with *known* parameters. After all, that would be a routine mechanical exercise. Rather, it reflects deeper uncertainty about the *fundamentals* of the currently active DGP itself. This uncertainty about fundamentals becomes even more salient when we think about questions related to a future DGP, such as: how would the US-China relationship evolve in the coming 10 years and how would it affect wages and capital returns in, say, the US? Or how would "artificial intelligence" affect the productivity of labor and capital in the coming 10 years? These questions can—and are—approached with available evidence from the present and past. However, as both popular discussion and the scientific literature show, the available evidence does not narrow down the set of a priori plausible answers significantly.[1] The conclusion from this is that there are circumstances in which the assumption that agents know the DGP is less plausible. This raises the question of how agents may make forward-looking decisions in the face of such uncertainty.

We dub a setting in which there is substantial uncertainty about the DGP one of *fundamental uncertainty* (FU). We mean a type of uncertainty that goes beyond the question of whether a particular data series is generated by an AR(2) or an AR(5) process. To further clarify our notion of FU, let us describe three—partly overlapping—constellations of a DGP that we associate with FU. First, the true DGP may be stable but relatively complex, in the statistical sense, such that a large number of (structural) parameters is needed to describe or predict the process. As a consequence, a correspondingly large amount and time range of data is needed to reliably estimate these parameters. Available real-world data may well fall short of this requirement. Second, the DGP may be simple in a local environment (corresponding to the recent past), but unstable globally; it may be very hard to estimate and predict the corresponding structural changes. This may result from ongoing institutional and technological innovation in vibrant liberal market economies (Binswanger and Oechslin 2021). Third and related to the first two points, available data may be compatible with multiple DGPs that coincide locally but may widely differ globally (cf. "observational equivalence", "indeterminacy").

Importantly, there is nothing inherently "philosophical" about FU as we understand it here. It is reminiscent of an experience frequently associated with empirical work, namely that there are instances where a given set of data can lead to widely differing conclusions depending on identifying assumptions. This phenomenon is not only encountered by empirical researchers but also by real-world economic agents when making sense of their current economic environment in order to make intelligent decisions.

Rational expectations rely on "knowing" the "correct" DGP. This is (approximately) feasible in an environment about which agents learn through frequent interaction under similar circumstances, assisted by timely feedback. As a consequence, agents act as if equipped with a large representative data set—resulting from exactly these interactions—that allows for reliable and conclusive parameter estimations for this environment. The associated learning process allows agents to behave as if they knew the true DGP. In this sense, agents "know", e.g., the physics-based DGP that applies to driving a car. This physical DGP is rather time-invariant, even if car-types and breaking systems change occasionally and some reactions have to be re-

[1] As examples in the domain of wages and labor income consider, for instance, Segal (2018), Cette et al. (2019), Acemoglu (2021), Acemoglu and Restrepo (2021), and Autor et al. (2021).

trained. However, for many real-world economic decisions, circumstances do not provide such favorable learning opportunities. Even academic economists engage in long-lasting debates, e.g., about possible structural changes in the labor share of GDP and—related— labor income processes at various skill levels (see references in Footnote 1). In the narrow sense that, for some economic decisions, it is not possible to pin down the underlying DGP to a narrow set of conclusive estimates, we may say that rational expectations are "infeasible" under FU.[2]

The question we pose in this paper is how economic agents can make forward-looking decisions under FU in a way that exhibits cognitive/psychological plausibility. Given the considerable challenge of this question, we scale it back to a narrower focus. We consider consumption choices of a representative agent who is confronted with an *exogenous* stochastic stream of labor income at the aggregate level of the economy. The agent has a goal of consumption smoothing as in standard models. However, the standard optimal solutions are not feasible since the agent does not know the DGP of the exogenous labor income and is aware of this ignorance. The only information the agent is equipped with are past realizations of labor income. However, the agent is unsure about the validity of extrapolating from this data into the indefinite future and is only prepared to use them for near-term forecasts. Moreover, even for near-term forecasts the agent does not consider the entire available history as relevant but considers certain episodes as particularly informative while others are not considered at all. This will give rise to making near-term forecasts by means of *historical analogies*. They reflect that the agent perceives history as generated by a potential mix of DGPs and is only willing to extrapolate from episodes that appear most similar to the current one. Technically, we implement this using the k-nearest-neighbors (knn) framework from statistical learning. Our assumption of forecasting based on similarity is motivated by the common observation that, e.g. in the news, sense making of current economic circumstances often occurs by comparison to historical reference episodes that appear similar to the current situation. For instance, pronounced downturns are often compared to 1929 or, more recently, to the great financial crisis; in the covid crisis there has been a frequent comparison to the supply shock of the 1970s (see Shiller 2019 for an in-depth discussion).

The key element of our model of forward-looking consumption choice is the separation of the "future" into two parts. A near-term future is associated with the time-increment from t to $t+1$.[3] The FU agent makes a forecast for labor income in $t+1$ based on similar patterns

[2] Compare this to empirical estimators that are commonly dubbed "infeasible" since certain parameters that are required as inputs are unknown under real-world circumstances.

[3] This time clock need not coincide narrowly with the frequency of time series data in empirical and simulated series. For instance, we may simulate quarterly consumption decisions, but the near-term future may comprise

in historical data. For the future from $t+1$ onward, the agent does not make any projections about the labor income process. This means that the value function that the FU agent uses for valuing resources available in $t+1$ is completely unspecific with respect to any DGP that may prevail from period $t+1$ on. This is the way in which the FU agent deals with the ignorance about future DGPs. The single defining property of the FU value function is increasing but diminishing marginal value of more resources. Intuitively, the FU value function indicates how the agent "generally" feels about ending up with, say, \$100,000 next period, in comparison to, \$90,000, \$110,000, ..., *whatever the future DGP of labor income may be.* The FU agent uses this value function for evaluating trade-offs between current consumption and future consumption *possibilities*. This evaluation is feasible exactly because it does not depend on the knowledge of the respective DGP. Overall, the behavior of the FU agent can be understood as a feasible version of rational expectations behavior. As we will argue in the main part of the paper, our approach is rooted in the recursive approach of dynamic programming and can be interpreted as a coarse "feasible" version of it that is based on learned estimations and approximations. We do, however, not provide any formal mechanisms of this learning process and leave it for future work.

In principle, the approach that we propose is in no way specific to a macroeconomic or a representative agent perspective. In essence, it may apply to any intertemporal decision at any aggregation level. However, we prefer to start with a macroeconomic perspective since this simplifies certain aspects of the analysis. For instance, this allows us to ignore life cycle patterns in income, or the impact of individually experienced events on individual belief formation. Rather, our focus is a collective aggregate element of sense making of current economic circumstances, and of consumption choices and consumption smoothing.

To test the explanatory power of our model, we run simulations based on aggregate macroeconomic data from the US. We first consider a stylized setting with only three income levels, derived from US GDP. The simple three-level setting serves as a laboratory environment to understand the working of our model and how it compares to a standard rational expectation model. We then proceed with a continuous income series estimated from US GDP. Our model captures well the strong comovement of consumption with income present in empirical data. We compare our predictions to a rational expectations benchmark model. The latter can also generate this strong comovement, but only for a relatively low discount rate. However, this leads, at the same time, to predicted asset holdings that show a negative correlation with empirical values. By contrast, our model explains both consumption and asset levels comparatively

an entire year. We will be more precise about this in the main analysis.

well. This is true in spite of the fact that our model does not exhibit an a priori higher degree of flexibility than the rational expectations benchmark model. A further negative aspect about the rational expectations benchmark is that its predictions become extremely sensitive to the discount factor as soon as average asset holdings increase beyond levels with binding borrowing constraints. On the negative side for our model, the predicted level of consumption smoothing is somewhat too low. A further interesting insight from fundamental uncertainty model is that the way in which agents form expectations may matter relatively little for predicted consumption and asset profiles.

The rest of the paper is structured as follows. Section 2 discusses related literature. Section 3 derives the FU model with the FU value function and forecasts based on historical analogies as key concepts. In Section 4 we provide a discussion of important aspects of the model in a question-and-answer format. In Section 5 we explain the empirical benchmarks that we use for simulations and for evaluating our model. The simulations are presented in Section 6. Section 7 concludes.

2 Related Literature

A seminal contribution to how real-world agents solve dynamic choice problems for which the optimal solution may be infeasible is the directed cognition model of Gabaix et al. (2006). In their model of a complex multi-step search task, the agent adopts a feasible approach by following an algorithm that replaces infeasible values by myopic approximations that avoid deep searches into a tree. Eventually, the search into the tree is stopped and the terminal node is evaluated with an unspecific fallback value.

Gabaix (2014) provides a static and Gabaix (2016) a dynamic model of how an agent manages to solve a complex decision task using "sparse optimization". In this approach, the agent has the notion of a default utility value resulting from not paying attention to a specific details of a decision task. The agent uses this default as the basis for mental simulations of whether it is worth to spend cognitive effort of taking a respective detail into consideration. This depends on the associated cognitive (search) cost and the expected utility benefit, where the latter is determined by the expected variability in utility outcomes from the respective factor of the decision problem. Our own approach can be seen as related in that our FU value function may be taken as a default. In other words, our FU agent does not pay attention to any element in the value function that relates to a future DGP since the cognitive costs are excessive.

Major contributions to decision making under uncertain DGPs have been made by Thomas Sargent and Lars Peter Hansen using the robust control framework (Hansen and Sargent 2008, 2010, 2022). In robust control, the agent doubts that the model of the fundamental DGPs underlying the economy may be misspecified. The agent therefore considers an entire range of models that are perturbations of a reference model. The agent then chooses actions that lead to good outcomes not only under the reference model but also under all perturbed models. In contrast to the robustness approach, the true DGP in our approach may be arbitrarily different from the model posited by the agent. They are only linked insofar as the agent observes data produced by past DGPs and uses this data for expectation formation. While robust control is very attractive from a normative point of view for rational policy decisions in the presence of model uncertainty, our approach is more geared towards a positive understanding of how agents make intertemporal choices under FU.

Woodford (2019) studies monetary policy in an equilibrium model of an economy with boundedly rational agents who lack the cognitive ability to calculate value functions that take into account an indefinite future. In Woodford's model, agents derive forward-looking behavior from forward planning, i.e. calculating forward into the event tree for a limited number of steps. The value of the stopping node in the tree, rather than being set to a default value, is estimated based on experiences collected in the history up to time t.

The key element that relates our approach to the above-mentioned studies is our FU value function. The second essential feature of our model is expectation formation for forecasting labor income in period $t + 1$. Here, our approach links to the concept of experience-based learning by Malmendier and Nagel (2016) and Malmendier and Shen (2018). Under experience-based learning an agent does not estimate the properties of a time series such as inflation or labor income based on an equal weighting of all available information. Rather, and mainly unconsciously so, the agent gives a higher weight to observations experienced during formative years at a young age. In our case, agents also give rather unequal weight to different historical episodes in that only episodes that appear similar to the current one are taken into account. Rather than individual forecasts at the micro level, we consider collective elements of forecasting concerning the economy-wide macro level. Expectation formation based on historical analogies links our approach to narrative economics (Shiller 2017, 2019). In fact, Shiller argues that historical analogies are one prominent way in which narratives appear and shape beliefs that are used for forecasting macroeconomic outcomes.

More generally, our study relates to the branch of statistical learning known as reinforcement learning which underlies dynamic algorithms of artificially intelligent behavior, e.g., in robotics or AlphaGo (see Sutton and Barto 2018). Unlike in economic models of rational expectations, when designing an AI system that is to make intelligent dynamic decisions, it cannot be assumed that the system already knows the optimal solution to a dynamic optimization problem. For many problems, including playing the game of Go and more complex control problems, the calculation of the optimal solution is infeasible. In a nutshell, reinforcement learning deals with how a learner can estimate a value function from data obtained in interaction with an environment, such that this value function leads to decisions that achieve good results. The derivation of our model borrows much from the estimation and approximation logic of reinforcement learning.

Finally, our approach is also related to the feasibility goals developed in Binswanger (2011, 2012). These are budget goals that an agent aims to achieve when it is too demanding to anticipate actual optimal future behavior, and hence a standard value function is again "infeasible". Specifying these budget goals still requires substantial cognitive skills and attention from the part of the agent. The approach put forward in this paper is therefore rather simpler.

Overall, the mentioned literature is mainly concerned with bounded rationality. Optimal solutions to dynamic problems are not available since agents lack the cognitive capacity to identify them. By contrast, our main theme is fundamental uncertainty. Identifying the appropriate DGP underlying current and future economic developments would still be challenging or even infeasible if mental calculation capacities were unlimited. It's the evidence on the current and future GDP that is the bottleneck. In reality, obviously, fundamental uncertainty about DGPs, bounded rationality, and forecast based on selective evidence are intermixed.

3 An FU Model of Intertemporal Consumption

3.1 Derivation of the FU Model

We consider an intertemporal consumption problem for a representative agent where budget and borrowing constraints are as follows:

$$\begin{aligned} c_{t+h} &\leq a_{t+h} + y_{t+h} \\ (a_{t+h} + y_{t+h} - c_{t+h})(1+r) &= a_{t+h+1} \\ a_{t+h} &\geq \bar{b}_{t+h}. \end{aligned} \tag{1}$$

In terms of notation, t denotes current time, and $h \geq 0$; c_{t+h} denotes consumption, y_{t+h} exogenous labor income, a_{t+h} assets, r returns on assets, and $\bar{b}_{t+h} \leq 0$ denotes a borrowing limit that we will later specify as a function of (exogenous) income. Labor income is thought to include transfers (such as income from Social Security), but we refer to it simply as labor income. It is this variable that is subject to FU. For simplicity, we treat returns on assets as deterministic to focus on a single source of (fundamental) uncertainty.

To start, it is helpful to consider the intertemporal consumption problem in a world where labor income is stochastic but drawn from a known process. The standard formulation of the decision problem is

$$\max_{(c_t, c_{t+1}, c_{t+2}, \ldots) \in \bar{\mathbb{B}}_t} \mathbb{E}_t \sum_{h=0}^{T-t} \beta^h u(c_{t+h}) \tag{2}$$

with $\bar{\mathbb{B}}_t$ referring to the sequence of budget and borrowing constraints in (1). Conceptually, solving (2) in the stated form is highly complex even if the agent knows the DGP of labor income, especially for large (or infinite) T and state spaces. It corresponds to a complete action plan for each possible state at each point in time. Graphically speaking, this means choosing an optimal consumption level at each node of a corresponding event tree, where the nodes corresponds to possible realizations of labor income. Even for numerical solutions of the standard version, the problem statement (2) is intractable, except for special cases. It is therefore rewritten in recursive form, as if the decision problem consisted of a sequence of simple decisions in a two-period world. With V denoting the value function, the above problem becomes then

$$
\begin{aligned}
V_t(a_t, y_t, \bar{b}_t) = \max_{c_t} \left\{ u(c_t) + \beta \mathbb{E}_t V_{t+1}(a_{t+1}, y_{t+1}, \bar{b}_{t+1}) \right\} \quad s.t. \\
c_t \leq a_t + y_t \\
(a_t + y_t - c_t)(1 + r) = a_{t+1} \\
a_t \geq \bar{b}_t
\end{aligned}
\tag{3}
$$

Note that for infinite-horizon decision problems, we have $V_{t+1} = V_t$ (provided the necessary technical conditions hold). This problem can be solved easily once the agent knows V.[4]

While—at the surface of it—the complexity of solving (2) has disappeared in (3), it is now hidden inside V_{t+1}, which represents the continuation value of ending up in the next period with assets a_{t+1}, labor income y_{t+1}, and a borrowing limit of $\bar{b}_{t+1}$. Importantly, V_{t+1} depends on a continuation with optimal decision making from $t + 1$ on. This means—and here comes a circularity (or rather the recursive element)—that problem (3) can only be solved in this simple form if its solution for future time periods is already known. In dynamic programming, for relatively small T and/or small state spaces, this can be addressed by solving the problem backward from the last period and recursively substituting the optimal continuation solution. In general, in dynamic programming with an infinite time horizon this circularity is solved by understanding the first line in problem (3) as a functional equation for V that, if everything is well-behaved, has a unique fixed point that can be found with the help of a computational algorithm such as value iteration. It is usually assumed that economic agents behave "as if" they knew this fixed point solution.

Even under FU, (3) is a natural starting point for intelligent forward-looking intertemporal behavior. This has been demonstrated in the literature on reinforcement learning (underlying dynamic AI and the successes of AlphaGo). The key idea is that the agent uses *estimations* of the complex composite term $\mathbb{E}_t V_{t+1}$ in (3). Very loosely, we may restate a "feasible" version of (3) as

$$
\max_{c_t} \left\{ u(c_t) + \beta \widehat{\mathbb{E}_t V_{t+1}(a_{t+1}, y_{t+1}, \bar{b}_{t+1})} \right\}. \tag{4}
$$

There are two components under the hat symbol in (4), an expectation and a value function component. The expectations component—$\mathbb{E}_t$—incorporates stochastic elements associated with the transition from t to $t + 1$. The second element—the value function component—

[4]As it is well known, there are cases where no knowledge of the value function is needed in order to solve this consumption problem as the solution can be inferred from the Euler equation. This only holds for special utility functions and the simplest versions of feasibility constraints.

is a utility-based measure of the value of ending up in a certain state in $t+1$ given by a constellation of assets, labor income, and a borrowing constraint. A priori, the true (but possibly "infeasible") value function depends not only on the utility function u but also on the properties of the labor income process $\{y_{t+h}\}$, $h \geq 1$. To see this, note that for a process with an overall (unconditional) higher probability mass on favorable outcomes of labor income y_{t+h}, V_{t+1} would take comparatively higher values for any given state in $t+1$. Further information that is "baked" into V_{t+1} relates to feasibility constraints.

We follow the literature on bounded rationality and reinforcement learning of replacing the backward logic of recursive substitution by a forward logic according to which $\mathbb{E}_t V_{t+1}$ is partly estimated using information that is currently available. We also follow the idea of the literature that, at some point, taking account of the further future occurs by using a default valuation that does not depend on any specific information extending beyond a certain point in time.[5] In our case, we assume that the agent uses past information for forecasting y_{t+1}, or—with a little abuse of notation—for an estimate $\widehat{\mathbb{E}}_t$. By contrast, the agent uses an unspecific default value function for $\widehat{V}_{t+1}$, denoted by V^{FU}. This default value function has the properties that it is strictly increasing in all of its arguments and that the marginal increases decline.

Our aim is to obtain an FU model that is as simple and tractable as possible, such that we can easily attribute simulation results to particular features of the model and we can even obtain a simple analytical expression for the optimal consumption choice. For this, we make three simplifying assumptions. Here, we simply state the assumption and explain their technical meaning. We discuss and justify them in more detail in Section 4. Our first assumption is that

$$V^{FU}(a_{t+1}, y_{t+1}, \bar{b}_{t+1}) = V^{FU}(a_{t+1} + y_{t+1} - \bar{b}_{t+1}). \tag{5}$$

Thus, the arguments are simply summed (recall that $\bar{b}_{t+1} \leq 0$). To understand this, consider the case that $a_t \geq 0$ and no future borrowing is possible, i.e. $\bar{b}_{t+1} = 0$. In this case, $a_{t+1} + y_{t+1}$ indicates available resources in $t+1$. Suppose now that $a_t \leq 0$, $a_{t+1} + y_{t+1} \leq 0$, i.e. in the absence of further borrowing opportunities in $t+1$, no positive consumption level would be feasible in $t+1$. However, if $a_{t+1} > \bar{b}_{t+1}$, the agent can still borrow up to an amount $a_{t+1} - \bar{b}_{t+1}$ in $t+1$ and so this adds to the feasibility set for future consumption. The FU value function directly gives a value for this feasibility set.

Our second assumption is that the agent relies on a certainty equivalence logic. Thus, the agent does not consider an expected value of V^{FU} for different possible values of y_{t+1}.

[5] See section 2 above.

Rather, the agent makes a single forecast for y_{t+1} and the FU value function depends only on this forecast. Denote the period-t point-forecast for y_{t+1} by $y^*_{t+1|t}$. Furthermore, since in our simulations the borrowing constraint will depend on labor income, denote the respective forecast for the borrowing constraint by $b^*_{t+1|t}$.

Third, for simplicity, we assume a logarithmic functional form for the FU value function, and also for $u(c)$. In sum, we thus have a feasible version of (4) that reads as

$$\max_{c_t} \; \log(c_t) + \beta \log\left(a_{t+1} + y^*_{t+1|t} - \bar{b}^*_{t+1|t}\right). \tag{6}$$

Given a forecast, this is an extremely simple decision model akin to a two-period consumption savings model or a two-period OLG model. The optimal consumption decision is given by

$$c_t = \begin{cases} \frac{1}{1+\beta}\left[a_t + y_t + \left(y^*_{t+1|t} - \bar{b}^*_{t+1|t}\right)/(1+r)\right] & \textit{if borrowing-constraint not binding} \\ a_t + y_t + \left(y^*_{t+1|t} - \bar{b}^*_{t+1|t}\right)/(1+r) & \textit{otherwise} \end{cases} \tag{7}$$

We directly see that more optimistic forecasts lead to higher consumption. We now consider forecasting, before we turn to a general discussion of the model elements.

3.2 Forecasts from Historical Analogies

Following the narrative economics approach (Shiller 2017, 2019), we assume that forecasts are built on historical analogies. The rationale is as follows. If the agent anticipates that the DGP may change in the future in unpredictable ways, then logical consistency suggests that also past events may have been generated by a DGP exhibiting changing behavior. One may imagine that different local versions of a DGP have been at work. Hence, the agent does take all past data as (equally) relevant for forecasting $y^*_{t+1|t}$, but only data from episodes that are "similar" to the current one. Put differently, the episode logic provides a sample selection procedure for the forecasting problem.[6] In what follows, it is important to keep in mind that we consider a representative agent. Thus, the expectation formation based on past episodes should be seen as a *collective* sense making process fed by numerous individual contributors.

Denote by $Y_t = (y_1, y_2, \ldots, y_t)$ the vector of all available historical data on labor income (including y_t). Denote by n^b the number of lagged periods that the agent considers as the

[6] It is an interesting question of how the performance of episode-based forecasts may compare to standard forecasts based on autoregressive or state-space models. Our current focus is on a forecasting model that we see as cognitively and psychologically plausible under FU.

"immediate history" of y_t, *including* a lag of 0. For instance, n^b may amount to 8 quarters. Let Y_t^b denote the subset of Y_t that contains only the data back to $t - n^b + 1$ and includes y_t. Let $x_{tj} := f_j(Y_t^b)$, $j = 1, 2, \ldots, J$, where $f_j(Y_t^b)$ denotes any suitable function. Finally, define $x_t := (x_{t1}, x_{t2}, \ldots, x_{tJ})$. In the language of statistical learning, x_t contains the features (or explanatory variables) for a prediction of y_{t+1} (the "target"). For our baseline simulations, we will use weighted identities and first differences as functions f_j, i.e.

$$x_t = (w_0 y_t, w_1 y_{t-1}, \ldots, w_{n^b-1} y_{t-n^b+1}, w_0 \Delta y_t, w_1 \Delta y_{t-1}, \ldots, w_{n^b-2} \Delta y_{t-n^b+2}), \tag{8}$$

where $\Delta y_t \equiv y_t - y_{t-1}$, and $w_l > 0$ denotes weights for lag l.[7]

Let $\|x_t - x_{t-l}\|$ denote the Euclidean distance between x_t and x_{t-l}, $n^b \leq l \leq t - n^b$ (where the date of the first available observation is $t = 1$). The most similar past period $t - l^*$ is given by

$$l_1^*(t) := t \ - \underset{l \in \{n^b, n^b+1, \ldots, t-n^b\}}{\operatorname{argmin}} \|x_t - x_{t-l}\| \tag{9}$$

Ranking lags according to $\|x_t, x_{t-l}\|$, we also obtain $l_2^*(t), l_3^*(t), \ldots$ Note that a period is only seen as similar if it falls outside the immediate history of y_t, which has length n^b (including y_t). We assume that the agent forecasts $y^*_{t+1|t}$ using the k-nearest-neighbor (knn) algorithm from statistical learning:

$$y^*_{t+1|t} = 1/k \sum_{m=1}^{k} y_{t-l_m^*(t)+1}. \tag{10}$$

In words, the agent takes the average of the lead values of the k past periods that have a history that makes them most similar to the current period.

It is worth noting that the knn prediction can be interpreted as a formalization of "retrieving a similar instance from memory". In statistical learning, knn it is characterized as "lazy learning". It does not require to first estimate parameters of a specific model. Rather, the agent simply looks at the history and forms expectations based on previous episodes that appear most similar. In other words, the agent "just" retrieves from memory what appears most similar. Since they are "model-free", forecasts based on knn are not constrained to be adaptive in the sense of "adaptive expectations" that are formed as some mechanical update based on previous changes or forecast errors. Rather, forecasts may jump from one period to the next. This can be seen as introducing an element of "animal spirits" into forecasting.[8]

[7]Specifically, for our baseline simulation we choose $n^b = 8$, $w_0 = w_1 = 2$, and $w_l = 1$ for $l > 1$. Thus, observations from the recent history enter the similarity calculations with a higher weight.

[8]Since it can "jump", knn learning can be very fast, provided a relevant similarity is visible to the agent. In this sense, in a more general setting, agents may quickly learn from policy makers that they may try to

An open question is how frequently the agent would make a new forecast. As a baseline, we assume that the agent does so in every period. We do so not because we think that this is particularly realistic. In reality, it is likely that agents are subject to limited attention and only change their forecast when this is triggered by a surprise in the form of a comparatively large forecasting error. This would lead to sticky expectations (Carroll et al. 2020). However, assuming a new forecast in every period allows us to keep the model more parsimonious as we do not need to specify the form of limited attention.

In the derivation above we have assumed that the forecast for labor income concerns the period $t+1$. For a simulation of the model and comparison to aggregate consumption data, we need—for empirical reasons—somewhat more flexibility. The macroeconomic series that we use for our simulation come with a quarterly frequency. This suggest that time moves in quarter steps. However, a priori, it is not plausible that an agent only considers the next quarter for a specific forecast $y^*_{t+1|t}$ and delegates everything beyond one quarter to the non-specific default valuation function V^{FU}. Rather, a plausible "takeover point" for the default may be one year. Formally, let the time horizon—expressed in baseline time steps such as quarters—when V^{FU} takes over denote by $t + n^f$. Thus, $n^f \geq 1$ is the maximum number of baseline periods that the agent can make a specific forecast. The forecast relies on similarities to past episodes and hence comes with an implicit assumption about a currently active DGP. A straightforward adaptation of (10) to this setting is then

$$\bar{y}^*_{t+n^f|t} = \frac{1}{n^f} \sum_{h=1}^{n^f} y^*_{t+h|t}. \tag{11}$$

With this, the agent makes a knn forecast not only for $t+1$, but up to a lead of $t+n^f$. For each of the episodes corresponding to a k value, the average of these forecasts is taken, and finally everything is averaged across the k values. The arguments of V^{FU} obtain a time stamp of $t+n^f$ (rather than $t+1$), and the "discount factor" β also has to be adjusted accordingly. Note that there is a degree of fuzziness here. The value of $\bar{y}^*_{t+n^f|t}$ does not change if the sequence of horizon-h forecasts is rearranged. This means that the agent's forecast reflect more of a fuzzy feeling about the near-term future rather than a forecast of a precise trajectory.

fool them, possibly more quickly than under a standard learning mechanism. This is of interest when linking forecasting to the Lucas Critique. In the current paper, there is no scope for policy makers to manipulate expectations since labor income is assumed to be exogenous, but this is a topic of interest for further research.

4 Discussion of the Model

For a better overview, we frame the discussion in a question-and-answer format.

What is the interpretation of V^{FU}? Let us assume $\bar{b}_{t+1} = 0$ for this discussion (see below for more on why the borrowing constraint is an argument of the value function). The true but "infeasible" value function V_{t+1} indicates the value of having resources of $a_{t+1} + y_{t+1}$ at disposal in $t + 1$. However, the true value function also depends on the DGP of labor income. For instance, the value of a high y_{t+1} depends on whether this increases the likelihood of persistently high labor income levels in the follow-up periods, or whether this would trigger mean reversion etc.

A simple way to think about V_{t+1} is to take it as an index. With proper rescaling, it may indicate how good it is to have, say, assets of \$100,000 and an income of \$150,000, on a scale from 0 to 10. However, as explained, the index value would depend in more or less subtle ways on the future DGP or labor income. If the agent is uncertain about the DGP and its future changes, then this dependence of the value function on the DGP cannot easily be expressed. In our model, we assume that the agent nevertheless finds it meaningful to ask: how good is it to have assets of \$100,000 and an income of \$150,000 in $t + 1$? The reason why this question may still be meaningful is that real-world agents have an idea of what this situation means, even if they may feel that they cannot make any sensible forecasts of what would happen in 5, 10, 15 years... This idea of what the situation means is based on learning from one's own experience, but there is almost certainly also a major degree of collective learning involved. We find it therefore plausible to assume that the FU agent in our model uses a value function that is unspecific with respect to a future DGP. V^{FU} represents this unspecific value function.

Possibly, in the real world, valuations of resource levels may still show traces of having been learned from specific past DGPs, or more generally from representations of the results of past DGPs in agents' memories. In fact, research by Malmendier and Nagel (2016) and Malmendier and Shen (2018) clearly suggests so, at least partially. To the degree that human learning is not excessively overfitting to past experiences, it is still plausible to assume that there is a component in these valuations that is, if not constant, only slow-moving. Here, we abstract from any historical dependency of V^{FU} and reduce our a priori assumptions to a minimum: a positive first and a negative second derivative, and a value of infinite marginal value at 0 (see below).

Why is borrowing part of the value function, and why does it enter additively? Suppose that the agent forecasts to end up in $t+1$ with $a_{t+1} = y^*_{t+1|t} = 0$. It may make a big difference whether the agent anticipates to be able to borrow and finance a positive consumption in $t+1$ or whether this is not possible. More generally, a low resource level in $t+1$ is less threatening if the agent has ample borrowing capacity compared to a situation with a lower borrowing capacity. Therefore, the borrowing limit must be an argument of the value function.

Since borrowed resources directly add to $a_{t+1} + y^*_{t+1|t}$, the additive form also makes sense. Suppose now that we assume $\lim_{t\to\infty} V^{FU\prime}(0) = \infty$. When the argument of V^{FU} is $a_{t+1} + y^*_{t+1|t} - \bar{b}^*_{t+1|t}$, this has the effect that the agent wants to avoid at all costs a situation with $a_{t+1} + y^*_{t+1|t} - \bar{b}^*_{t+1|t} = 0$, as it should be. If the borrowing part were missing in the argument, then it would mean that the agent avoids at all costs a situation with $a_{t+1} + y^*_{t+1|t} = 0$. This would imply that the agent would avoid at any costs that assets become lower than forecasted income, which seems somewhat arbitrary.

Why assuming a logarithmic functional form for V^{FU}? As discussed in the previous paragraph, it makes sense to assume $\lim_{t\to\infty} V^{FU\prime}(0) = \infty$. This means that the agent avoids at all costs a situation with a forecast of zero future resources. Apart from this, we want to impose that the marginal value of resources is positive but decreasing. The logarithm is the simplest and most tractable functional form that fulfills all these assumptions. Moreover, there is no extra parameter to be determined. However, as an extension, a CRRA functional form, or any other tractable functional form with the above properties would be a valid candidate for V^{FU} as well. With CRRA one may fine-tune the degree of consumption smoothing.

Why certainty equivalence in V^{FU} instead of $E_t V^{FU}$? Whether it is in official publications of forecasts for economic indicators or discussions in the news, we almost always find forecasts in the form of a single number $y^*_{t+1|t}$ rather than a distribution expressed with probabilities $\hat{p}(y_{t+1|t})$. In line with this observation, it is well documented that many individuals have difficulties with probabilistic forecasts (see Binswanger and Salm 2017 for a discussion). Finally, it is straightforward to model expectation formation and forecasting with historical analogies ("narratives") using a point forecast derived via knn. In a probabilistic model, the agent would first have to estimate the relevant scenarios and then the associated probabilities. It is not as straightforward to posit a parsimonious model of how this is done in combination with a similarity-based reading of past data. It is possible to extend the knn framework in this direction, but it requires several additional assumptions and is thus less parsimonious.

Why does the FU value function already "take over" in $t+1$ Suppose that an unspecific default version of the value function would only apply from some period $t+h^*$ on, with $h^*>1$, and that this would be explicitly modeled. This would imply an explicit forward search into the event tree from the time-t node onward until reaching a depth of h^*. Clearly, this would make the model less parsimonious and tractable. Also, the underlying planning processes posited by the model would be cognitively more demanding. As a consequence, our simple version where the unspecific default value function V^{FU} already applies from $t+1$ on seems a natural starting point for a model of forward-looking behavior under FU.

Still, it may be plausible that agents are in a position to make specific forecasts—at least on an intuitive level (see below)—ahead of what would literally count as $t+1$ according to the model's baseline time frequency (quarters, years). For this situation, we find it conceptually simpler to adopt the trick expressed in the forecasting formula (11), which we understand as a rough approximation to a somewhat deeper search into the event tree. It would be possible to refine (11) by taking into account the sequence of forecasts over horizons h, e.g. with using a weighting scheme. The disadvantage of this is that more parameters are introduced.

Is there learning? As we have argued above, V^{FU} can be seen as a result of individual and collective learning. We assume that this learning has led to the following three features: (1) additional resources lead to a higher value; (2) the additional value of more resources decreases; (3) a situation with zero resources (taking into account any borrowing options) are to be avoided at all costs. All three assumptions are natural in economics. The specifics of the underlying learning process may be tied to individuals' past experiences (Malmendier and Nagel 2016; Malmendier and Shen 2018). Here, we do not attempt to model this underlying learning process. Rather, we take V^{FU} as exogenously given and sufficiently slow moving, such that treating it as given is relatively innocuous.

Is there learning associated with forecasting in our model? While knn is a statistical learning model, in our context it may be more suitable to understand it as a model of memory retrieval. The agent does not learn a model that improves over time with more data. Rather, the agent's memory grows with more historical experience. Hence, forecasting does change over time in the sense that the agent has seen more reference episodes that may serve as an improved basis for further forecasts. Intuitively, the great financial crisis has become a reference period in the post-2007 world, but it had not been in the pre-2007 world. Clearly, a richer experience set need not necessarily improve forecasts. It may just increase the scope for overfitting. However, this paper is not about optimal forecasting but about developing a stylized positive model of forecasting of real-world agents that leads us to better understand empirical consumption choices.

Isn't it cognitively demanding to form forecasts based on knn? There is an "as if" logic behind the assumption of similarity-based forecasting via historical narratives. We do not claim that every single individual in an economy engages in a knn-like search into the past to identify periods that are particularly similar to the current one. Rather, we imagine that this is associated with a "collective mind". As documented in Shiller (2019), evidence suggest that economic circumstances are often assessed by means of similarity to the past. While it is sometimes possible to track the origins of a corresponding historical narrative, more often it is not, as these narratives are collective phenomena with many minds being involved into its emergence and spread. The knn-based forecasting in our model is a shortcut for this collective process.

Aren't the knn-based forecast very flexible, so we can generate any expectations? Compared to all the elements that may affect beliefs and expectations in the real world, the forecasting that we model is very restrictive. Our representative agent looks at one single aggregate historical time series. Variables like unemployment, inflation, financial markets, political events etc. cannot influence forecasts (except to the degree that they appear predictable from past income realizations). The only information used for calculating similarities are past levels and first differences of income. These are the simplest possible functions of past income values. The two most recent levels and differences obtain a higher weight such that periods in the past classified as similar have a high likelihood of sharing similar levels and recent directions of moving (up, down, or sidewise) with the current period.

Our approach could also be applied to expectations on an individual microeconomic level. In this case, a restriction of expectation formation from past individual income realizations may make less sense. Therefore, micro models would need to be somewhat richer. There, other criteria can be used to keep discipline on degrees of freedom, such as available evidence form micro data.

Is it all about the past? What about the great moderation? Our knn-based model cannot explain any forecasts that have a flavor of "this time is different". For instance, in our model, forecasts in the great moderation cannot be based on a narrative that central banking has become so much better and structured products would diversify away so much risk that business cycles and crises are a matter of the past. Rather, under knn, optimism during the great moderation is driven by a reference to past booms with a longer duration, and eventually by the a reference to the great moderation itself as it starts to represent a "new normal"—through the lens of the knn model.

Could the model get extended to incorporate life cycle patterns? Apart from a potentially unstable DGP for labor income, our model depicts a setting where circumstances stay identical. Most importantly, there are no life cycle patterns. However, even if the DGP or labor income is subject to FU, this does not necessarily mean that everything concerning the future is subject to FU. Even if it is hard to predict the wage level for a particular skills group in 10 years, individuals currently at the age of 45 may not feel very uncertain about the number of children they may still have after the age of 50 (probably few) and about retiring at an age close to 70 (at least if they are white-collar employees). In other words, some elements of the future are far more predictable than others, or at least they may appear more predictable to decision making agents. These (maybe only seemingly) predictable elements can be "baked" into V^{FU}. For instance, if our approach is translated to an OLG or microeconomic setting, then the value function may become also a function of age and it could feature a wealth level that is deemed as necessary for financing retirement and that is subtracted from otherwise available resources. Agents may be subject to limited attention and certain aspects of the environment may only appear gradually, e.g. retirement saving may only be a concern at a higher age. As attention may play an important role here, we see it as promising to explore how our approach could be combined with the sparsity/limited attention approach by Gabaix (2016).

5 Empirical Benchmarks

To explore whether the predictions of the FU model are compatible with empirical patterns in real aggregate consumption and net worth series, we calibrate it to the US economy. We do so in two ways. First, we derive simulations for a highly stylized income process with three income levels, based on aggregate US income data. We compare them to a rational expectations benchmark model. The simple income process allows for crucial insights into the mechanics of the FU model and key differences to a benchmark rational expectations (RE) model. Second, we run simulations for a continuous estimate of de-trended US aggregate income and compare them to empirical data. In both scenarios, the FU agent does not have any model of the income process in mind. For the rational expectations benchmark, we will assume that income process is Markovian.

There are several possible empirical counterparts for labor income in our model. As a baseline, we simply take quarterly real US GDP. In this approach, we imagine that labor income is a fixed proportion of GDP. However, with the procedure described below, multiplying GDP with a constant factor does not have any effect on the resulting series, therefore, we can directly take GDP as our raw measure of labor income. For the same reason, we also ignore taxes. We use a smoothing spline with 3 knots to estimate a trend for the logarithm of GDP. This provides a parsimonious and relatively "model-free" way to take into account that trends have changed due to productivity slowdowns. We then convert the trend back to levels, and calculate the difference between GDP and the estimated trend; we refer to these differences as residuals. We divide these residuals by the trend, such that the resulting series has a mean of 1. If any particular value of the series takes on the value 1, this means equality to the trend. The resulting continuous de-trended labor income series is shown as the black solid line in Figure 1.[9]

To approximate a three-level income process, we first specify a low, middle, and high income level. These are set to the three quartiles of the continuous series, respectively. Each value of the continuous series is then mapped to the one of the three levels that is closest to the original value. The result is presented as the discrete blue solid line in Figure 1. We also estimate the corresponding transition matrix using relative counts of transitions in the three-level series. The result is given by

$$\widehat{\boldsymbol{M}}_3 = \begin{bmatrix} 0.89 & 0.20 & 0.01 \\ 0.11 & 0.70 & 0.06 \\ 0.00 & 0.11 & 0.94 \end{bmatrix},$$

which shows a substantial degree of persistence.

For our estimate of the empirical counterpart of consumption we use quarterly personal consumption data from FRED.[10] The series is de-trended by taking residuals from a regression of the log of the original consumption series on the log trend for GDP, estimated with smoothing splines as described above. The fitted values from this regression provide the log-trend of

[9] A very similar picture is obtained if we do not use (a fixed proportion of) GDP as income measure but calculate labor income plus transfers using the respective series from FRED.

[10] We use the PCECC96 series.

Figure 1: Empirical income and consumption series

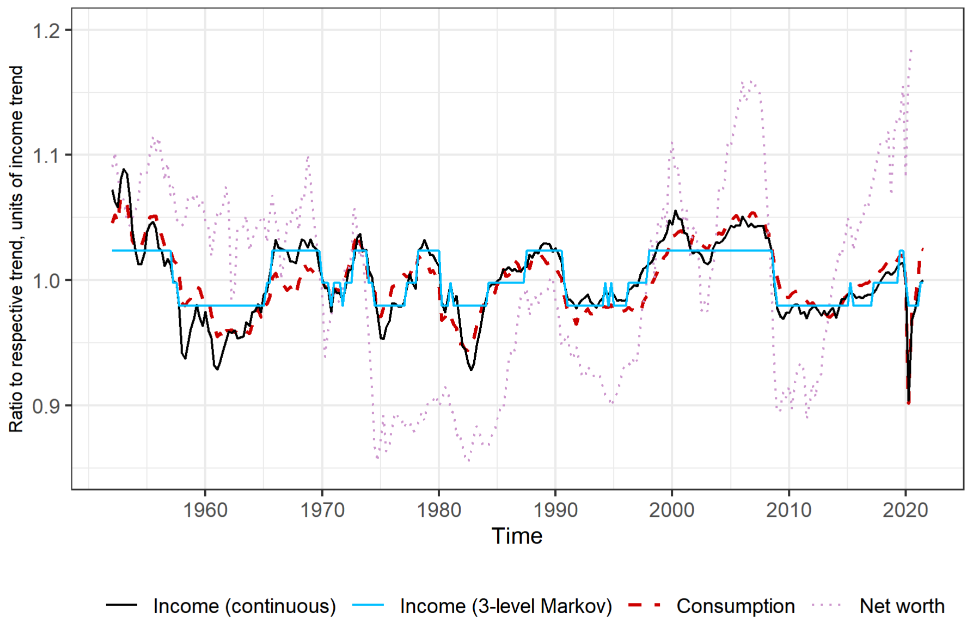

NOTE: The last four data points for net worth fall outside the range of the figure and have been removed for better readibility.

Figure 2: GDP, scaled GDP, and consumption: deviations from trends

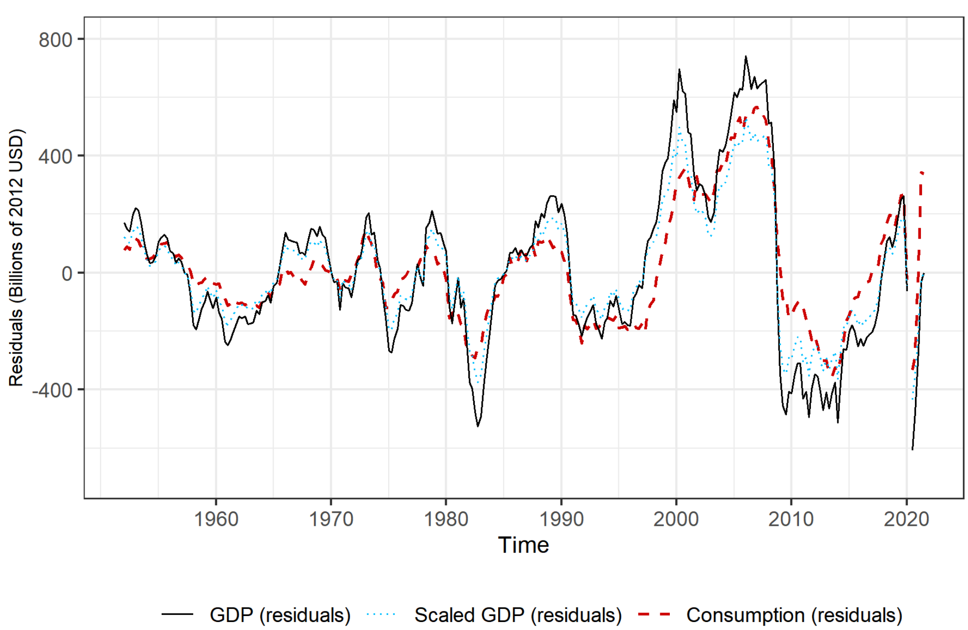

NOTE: The figure shows deviations from trends in levels (billions of USD). Data points for all series for the second quarter of 2020 fall outside the range of the figure and have been removed for better readibility.

consumption. The residuals are converted back to levels. We then divide them by the level-trend of income such that the resulting series is measured in the same units as income. Finally, we shift the series up by one unit such that it has the same mean as income. The resulting consumption series is shown by the red dashed line in 1.

The focus on consumption residuals and the normalization of their mean to one means that we can only study patterns in the consumption series relative to movements in income and the consumption series itself, not patterns related to the overall level of consumption relative to income or savings. This is consistent with the focus of this paper. The empirical mean of the raw values of personal consumption, as a share of labor income plus transfers, amounts to about .90, i.e there is a savings rate of 10 percent out of this measure of labor plus transfer income. Neither the FU model nor our benchmark rational expectations model (see Section 6) are tuned to explain a sustained savings rate out of labor income that differs substantially from 0. After all, the only savings motive present in the models we consider is consumption smoothing. Arguably, a major part of a sustained higher saving rate in the empirical data results from life cycle saving and from the desire to increase consumption in the future, or from running a business etc. These motives are absent in the models that we consider.

A further variable of interest is an empirical measure for assets/wealth, which results from accumulated savings. For this, we use the broad net worth measures from FRED.[11] Net worth has a trend that is substantially steeper than the trend in GDP and personal consumption. Again, neither our FU model nor our rational expectations benchmark features any elements that can explain this trend (which may relate, e.g., to high returns on capital, or increases in income and wealth equality). We therefore view it as legitimate to focus on patterns of net worth around its trend, even if that trend is different from the trend in the income and consumption series. This allows us to obtain at least crude evidence on how predictions of the FU model and an RE benchmark regarding asset levels relate to empirical movements in assets. With this in mind, we proceed with exactly the same steps as in the case of consumption. The resulting net worth series is shown by the dotted light-purple line in Figure 1.

Figure 1 directly shows the main empirical pattern: consumption is somewhat smoother than income, but the comovement with income is very strong.[12] The correlation coefficient for

[11]The measure of net worth that comes closest to our model is net worth for households (BOGZ1FL192090005Q). However, this is only available from the fourth quarter of 1987 on. An alternative measure, net worth of households and non-profit organizations (TNWBSHNO) is available for the entire sample period. We estimate the differences in means for the period where both measures are available and then subtract this difference from the series including non-profit organizations. For the conversion from nominal to real values we use the GDP deflator.

[12]This also holds when we use more direct measures for labor income and transfers.

income and consumption amounts to .88. To understand what this tight comovement precisely means, it is important to be careful with the units of the two series. For both series, the units are the time-t values of the income trend (in levels), and both series are based on deviations from the trend. An income value of 1.02 at time t means that the deviation of income from its trend at that time (in billions of US dollars) amounts to 2 percentage points of the value of the income trend. A consumption value of 1.02 means that the residual of consumption from its *own* trend amounts to 2 percentage points of the value of the income trend at time t.

To put this picture further into perspective, Figure 2 shows deviations from trends for income and consumption in *absolute* levels, measured in billions of real-2012 USD. The black solid line shows residuals for GDP. The blue dotted line shows residuals for GDP, multiplied by a scaling factor. This factor is equal to the mean of the ratio of labor income plus transfers to GDP. This is thus a rough measure of labor plus transfer income, expressed as deviation from the trend. The dashed line shows consumption, again in the form of deviations from the trend. Note that the swings in all series get larger when moving to the right. This is because—unlike in Figure 1—the series are not normalized by a trend, thus deviations from trends grow over time. Figure 2 confirms that consumption is only very moderately smoother than our measure of labor plus transfer income. However, it is significantly smoother than GDP.[13]

6 Simulations

6.1 Simulations for Three-Level Income Process

We first run simulations for the FU model for the three-level income series discussed in the previous section (see the blue line in Figure 1). Importantly, the FU agent does not estimate a model of this process, e.g. a transition matrix, but makes "model-free" forecasts using knn predictions that rely on similarities to past periods. For our simulations, we assume $k = 3$, a backward window of $n^b = 8$ periods for assessing similarities with past periods, and forecasts extend $n^f = 4$ quarters into the future. The similarities in the knn model are calculated using the Euclidean distance for the present, and past $n^b - 1$ levels and $n^b - 2$ first differences.[14] The current values and first lags of both enter the Euclidean distance with a weight of 2, the rest with a weight of 1. The higher weight has a relatively strong influence on which past periods are selected as most similar to the present one. It has the effect that the periods that form

[13]Note that, for improved readability, the values of all series in the second quarter of 2020 fall outside the range of the figure.

[14]With these numbers, the total elements of Y_t^b add up to exactly n^b.

Figure 3: Most similar past periods for forecasting

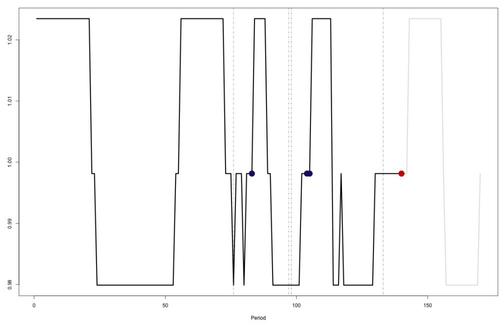

(a) Period 140

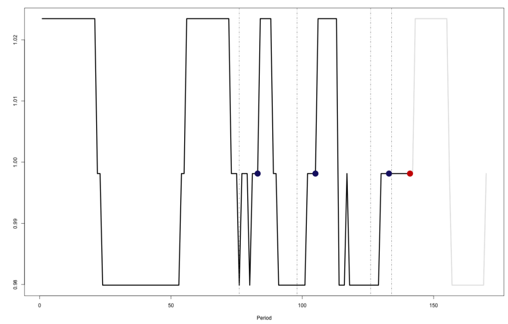

(b) Period 141

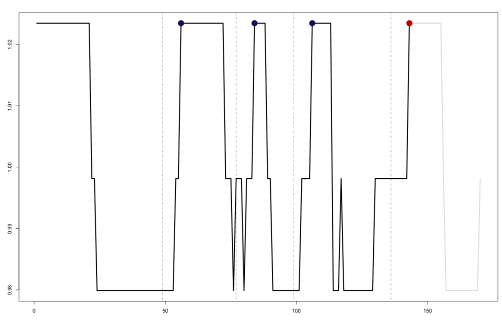

(c) Period 143

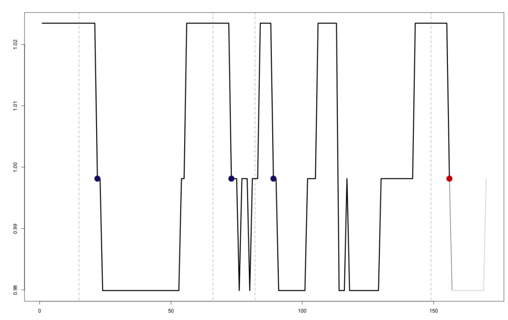

(d) Period 156

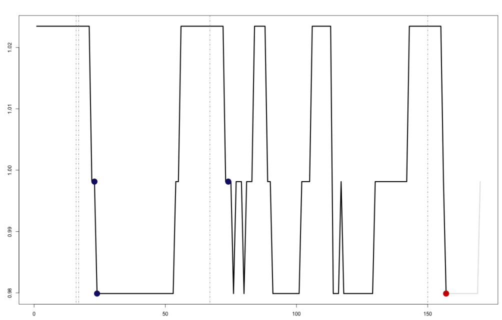

(e) Period 157

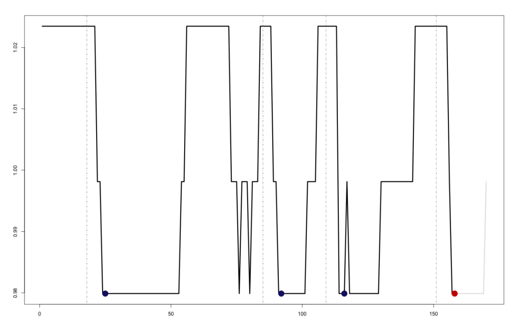

(f) Period 158

NOTE: The figures show most similar past periods for selected current periods. The main line shows the simplified three-level income series discussed in Section 5 and appearing in Figure 1. On the horizontal axes, time marks have been replaced with generic time count labels. In each subfigure, the red dot shows the current period, the blue dots indicate the three most similar periods from the past. The dashed vertical lines indicate the limits of windows into the past of length n^b (including the current period), starting at each respective colored dot. The true values of yet unrealized future income values are indicated by a gray color of the main income line.

the “historical analogy” share highly similar movements in the most recent two quarters of their respective histories. From a psychological point of view, this seems to be important for a perception of similarity.[15] For asset returns, we set $r = 0.02$.

Figure 3 sheds light on the agent’s assessment of which past periods are seen as similar to the present. For easier readability, the original time marks have been replaced with generic period labels. The red dot in each subfigure shows the respective current period. The blue dots show the three most similar periods in the past. The dashed vertical lines indicate the end of a backward window of n^b periods that is used for calculating the Euclidean distances. The blue dots always lie to the left of the right-most dashed vertical line that indicates the end of the backward window of the current period. Otherwise, similar periods can be arbitrarily close or apart from each other. Panel (a) shows how, in period 140, two of the most similar periods are immediate neighbors on the time axis. All periods similar to period 140 have a recent past with the lowest income realization, followed by an upward movement, and at least one period with a middle income level.

The panels in Figure 3 show some snapshots of how the income history unfolds. Incomes in the future of the respective red dots are shown as a gray line. From panels (a) to (f), the income line gradually blackens. The sequence shows how the past periods seen as most similar to the present often share the same income level and similar movements. The agent derives a forecast for y_{t+h} by averaging the lead values h periods to the right of the blue dots.

Figure 4 shows the results for consumption and assets for the FU model for various values of β, ranging between 0.5 and 3 (see legend). The profile labeled “Const. forecast ($\beta = 2$)” provides a comparison benchmark in which forecasts are not based on knn but are simply constant at the mean of the income series (and $\beta = 2$). The horizontal axis keeps year labels, as a reminder that the income series ultimately derives from US GDP. Panel (a) shows the results for consumption. The first striking observation is that all series strongly move together with income. Second, a low but economically significant level of consumption smoothing is also clearly visible. For $\beta = 0.5$, the standard deviation of the consumption series is 0.080. With a higher β, it decreases to 0.020 for $\beta = 2$. For β values exceeding 2.5, the standard deviations again increases slightly to 0.040 at $\beta = 4$ (not shown in the figure). Third, a higher β is associated with a slightly higher mean consumption—the means range from 1.002 for $\beta = .5$ to 1.006 for $\beta = 4$. Fourth, spikes in income forecasts—shown as the gray dotted line—trigger

[15] Readers who feel uncomfortable about this weighting scheme, or knn forecasts in general, potentially requiring too many parameters, may be more interested in an extremely simple alternative that we dub “constant forecast”. According to this alternative, the agent constantly forecasts an income level of 1. Calibrations for this benchmark are shown below.

Figure 4: Consumption and Assets for the FU model with 3 income levels income

(a) Consumption

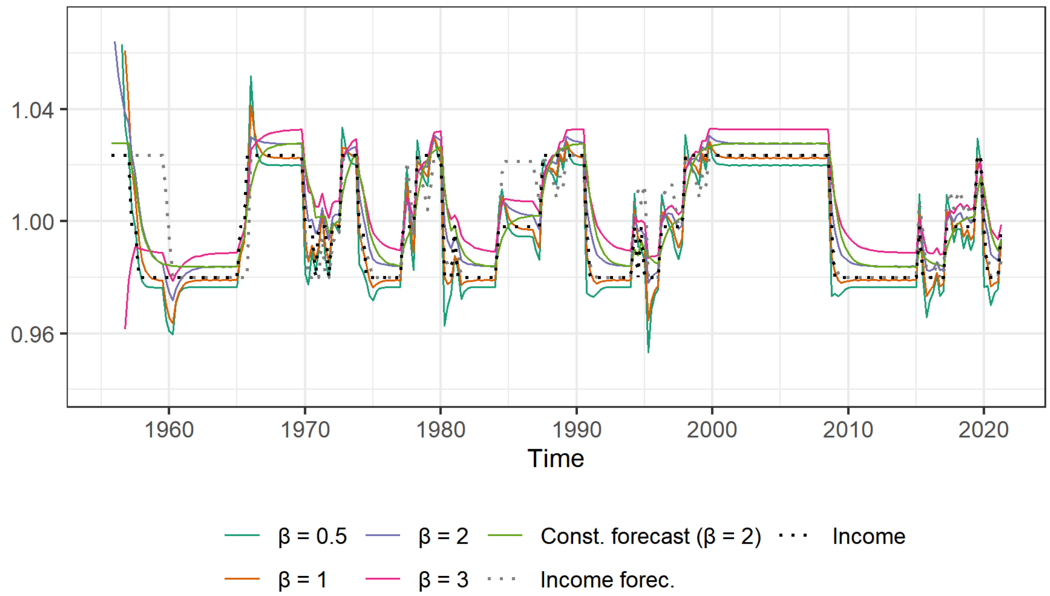

(b) Assets

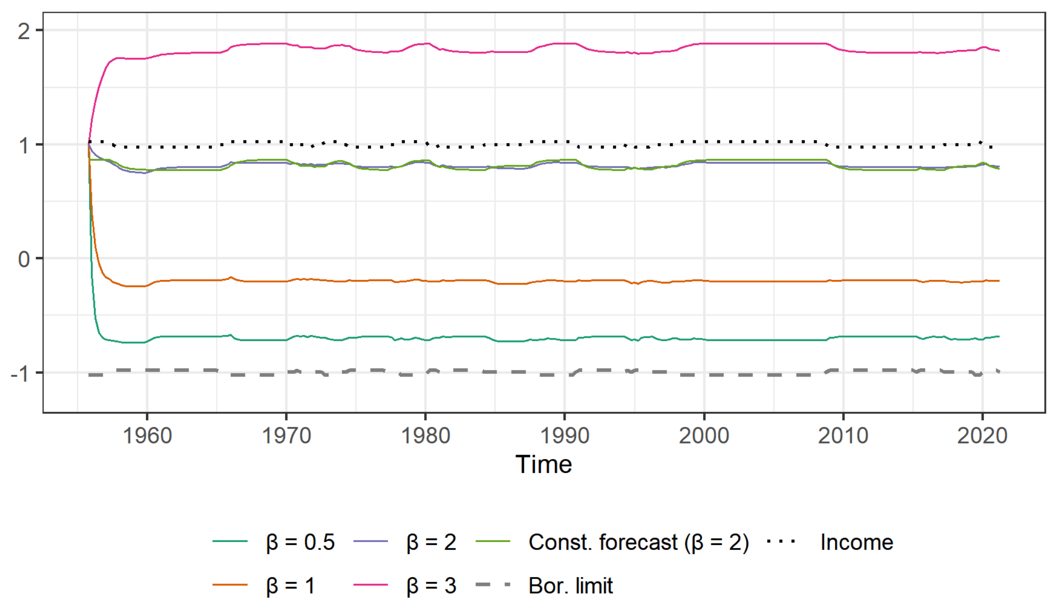

NOTE: The gray dotted line in Panel (a) shows the forecast $\bar{y}^*_{t+n^f|t}$. The parameter values are $n^b = 8$, $k = 3$, $n^f = 4$, $r = 0.02$. Similarity for knn forecasts is calculated with Euclidean distance for past $n^b - 1$ levels and $n^b - 2$ first differences. The current values and first lags of both enter the Euclidean distance with a weight of 2, the rest with a weight of 1. The line labeled "Const forecast ($\beta = 2$)" shows consumption for a constant income forecast equal to 1 (i.e. the mean). Initial assets are set to 1. The agent can borrow up to 100 percent of current income and anticipates to be able to borrow up to an additional 20 percent of $\bar{y}^*_{t+n^f|t}$ in the next period.

spikes in consumption. This becomes visible, for instance, after 1985, and before 2000. This shows that "expectations shocks" can trigger a consumption response. The constant-forecast version of the model does not share these spikes. It's standard deviation for the consumption series amounts to 0.017 and is thus the lowest of all series.

How does consumption smoothing occur in this simple model where the agent's expectations beyond one year are completely unspecific and the agent has no model of the DGP for income? The agent balances the marginal value of consumption in t with the marginal value of having more resources available in $t + 1$. The latter comes from the FU value function. For the simulations in Figure 4, initial assets are set to 1. In the first period, the forecast is a high level of income. Take a series with $\beta \leq 2$. Given optimistic expectations, the marginal value of future resources is relatively low and the marginal value of current consumption relatively high. Hence, consumption is relatively high. In the next period, wealth is already lower (see Panel (b)), but the constellation is still very similar. When the effective income drops in 1957, the agent first still forecasts a high income for a few quarters, so this represents a surprise series of negative income shocks. Based on the (overly) optimistic forecast, the agent chooses a consumption level that is relatively high. Hence, assets get depleted and the marginal value of future resources starts to increase and so consumption decreases. This process is gradual. Immediately before 1960, consumption is lower than income such that there are net savings out of labor income. This then leads to an increase in assets, such that the marginal value of future resources decreases. There is thus an ongoing rebalancing of marginal values of current consumption against the marginal value of future resources, with income shocks as external actors that impact this rebalancing process.

Panel (b) of Figure 4 shows the corresponding movements in assets. The most striking pattern is that different values of β lead to more or less parallel shifts in assets, after an adjustment from the initial asset level of 1. It is only for $\beta < 1$ that assets are persistently negative. At $\beta = 0.5$, they do not yet reach the borrowing limit. Since, under FU, it would be difficult to express borrowing limits in terms of the present value of "lifetime" income, they are expressed as ratios of y_t and $\bar{y}^*_{t+n^f|t}$; in the shown simulations, we set these ratios to 100 and 20 percent, respectively. To be precise, the agent anticipates to be able to borrow up to 20 percent of $y^*_{t+n^f|t}$ in the next period.

We now compare the simulation results for the FU model to a rational expectations (RE) benchmark. The decision problem by the RE agent is given by (2), with the budget constraint specified by (1). We assume $u(c_{t+h}) = \log(c_{t+h})$ and set T to infinity. The RE agent assumes

Figure 5: Consumption and Assets for the RE model with 3 income levels

(a) Consumption

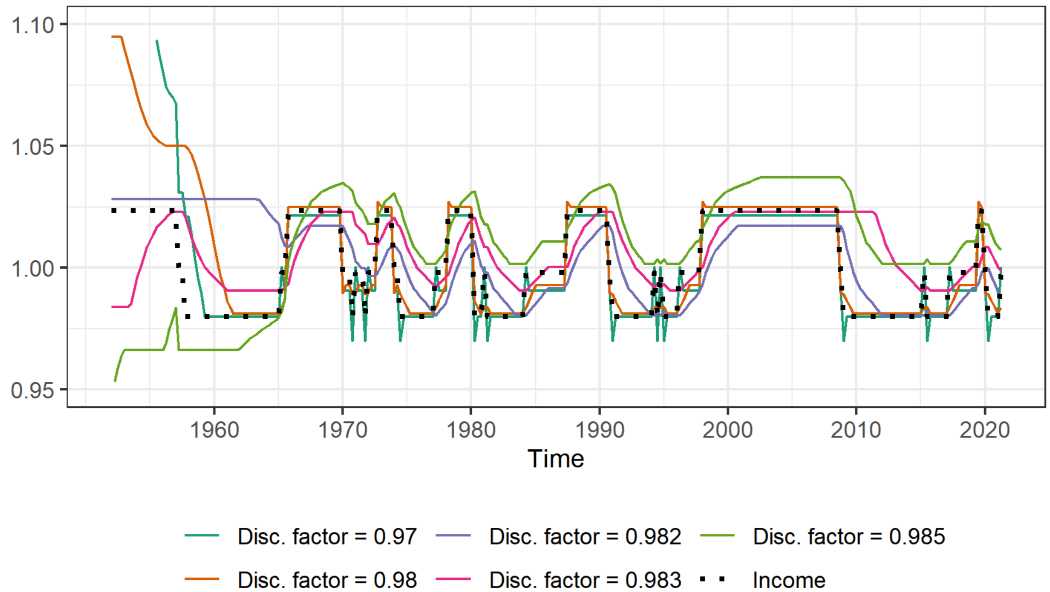

(b) Assets

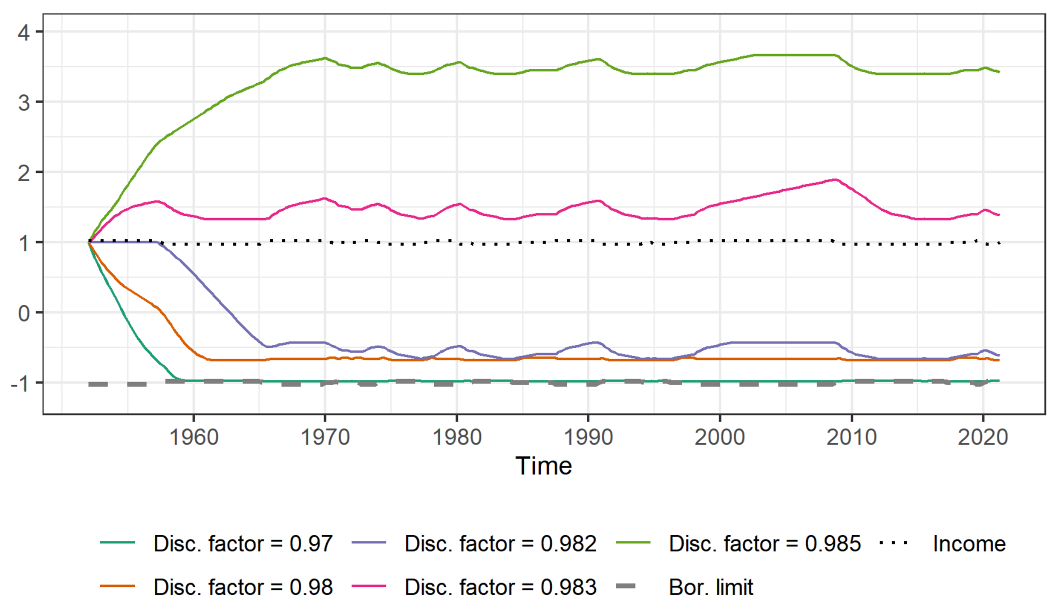

NOTE: The figures shows consumption and asset profiles for a rational expectations model with logarithmic utility function. The agent assumes that income is generated by a Markov process with transition matrix given by (5). The agent's borrowing constraint is 100 percent of current income. Initial assets are set to 1.

that the income process is generated by a Markov process with the transition matrix given in (5).[16] To distinguish the discount factor from the β parameter of the FU case, we refer to the former simply as such, i.e. "discount factor". The results for different discount factors are shown in Figure 5.

The most striking observations are the following. For discount factors of 0.97 and lower, assets reach the borrowing constraint. Hence, the movements in consumption follow those in income. The level difference is due to interest payments and changes in borrowing constraints triggered by income changes (recall that the borrowing limit is tied to current income). When the discount factor increases beyond 0.97, considerable smoothing appears very quickly. That consumption profiles still exhibits swings is due to the fact that assets are still negative and relatively close to the borrowing constraint. Furthermore, income changes show a high degree of persistence, as seen from the large diagonal entries in the transition matrix in (5). Third, and maybe most striking, assets are extremely sensitive with respect to the discount factor. Up to a value of 0.982, they are consistently negative. When the discount factor is increased by a tiny 0.001 to 0.983, assets mainly fluctuate around a value of 1.5. When beta increases to 0.985, they then fluctuate around a level of 3.5. For still higher discount factors, asset holdings quickly explode to levels around 30 and more. Thus, within a very narrow range of discount factor values, the qualitative as well as quantitative nature of the predictions of the RE model change dramatically.

6.2 Simulations for Continuous Income Process

Figure 6 shows simulation results for the continuous income series for the FU model. We show model predictions for a range of β values and similarity-based forecasts, as well as predictions for constant income forecasts equal to the mean and a $\beta = 2$. The thick yellow line shows the empirical consumption series (identical to the one shown in Figure 1).

Consumption profiles are shown in Panel (a). All FU consumption profiles exhibit a high correlation with income. For $\beta = 0.5$, consumption strongly reacts to changes in income forecasts since asset levels are relatively close to the borrowing limits (see Panel (b)). For higher β values, consumption profiles are slightly smoother than income. A main insight from

[16] A discussion of whether the Markov model is sensible or "rational", and whether knn forecasts may perform better than the ones based on the Markov model is not yet part of this version of the paper but planned for a next one. For the discrete income series that we consider here, it actually is true that there are only three income levels. However, the DGP may not be Markov. Our RE agent just assumes that the process is Markov and uses the estimated transition matrix for forming conditional expectations. However, the series has not been simulated by random draws according to the transition matrix and it is conceivable that the Markov model is misspecified.

Figure 6: Consumption and Assets for the FU model with continuous income

(a) Consumption

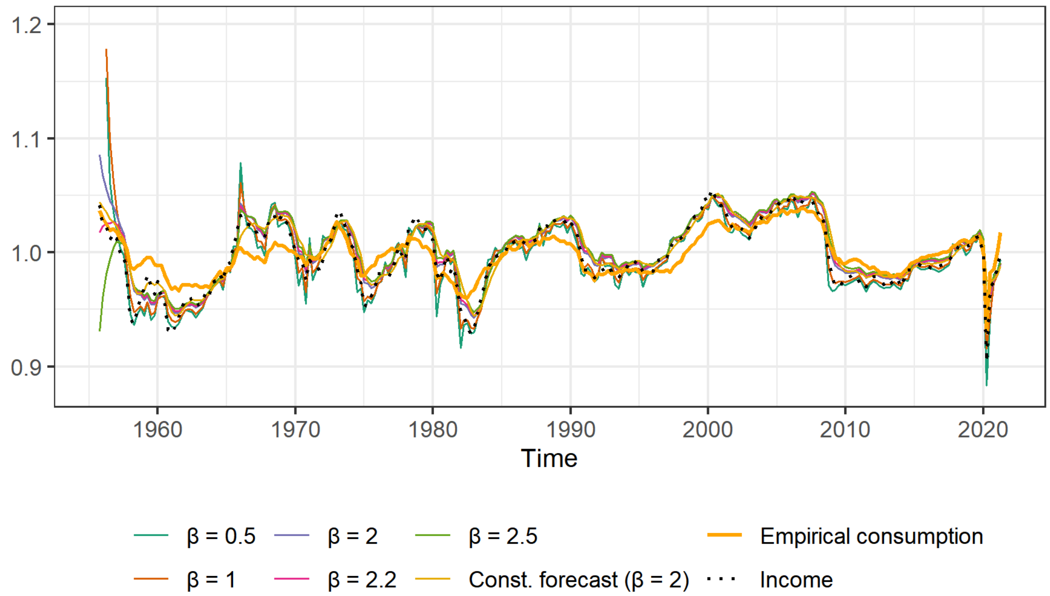

(b) Assets

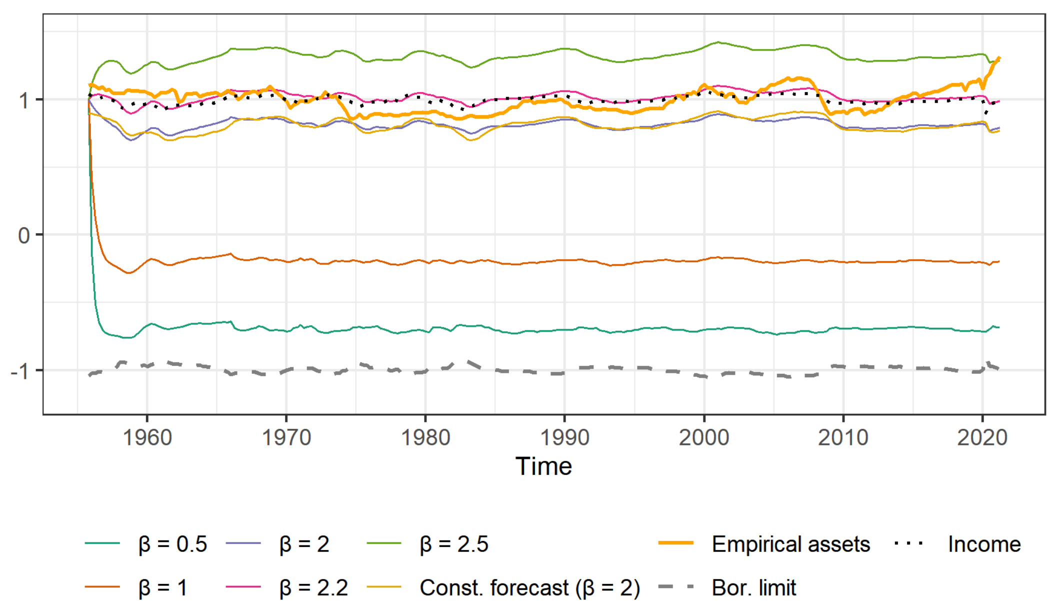

NOTE: The parameter values are $n^b = 8$, $k = 3$, $n^f = 4$, $r = 0.02$. Similarity for knn forecasts is calculated with Euclidean distance for past $n^b - 1$ levels and $n^b - 2$ first differences. The current values and first lags of both enter the Euclidean distance with a weight of 2, the rest with a weight of 1. The agent can borrow up to 100 percent of current income and anticipates to be able to borrow an additional 20 percent against $\bar{y}^*_{t+n^f|t}$ in the next period.

Figure 7: Consumption and Assets under RE

(a) Consumption

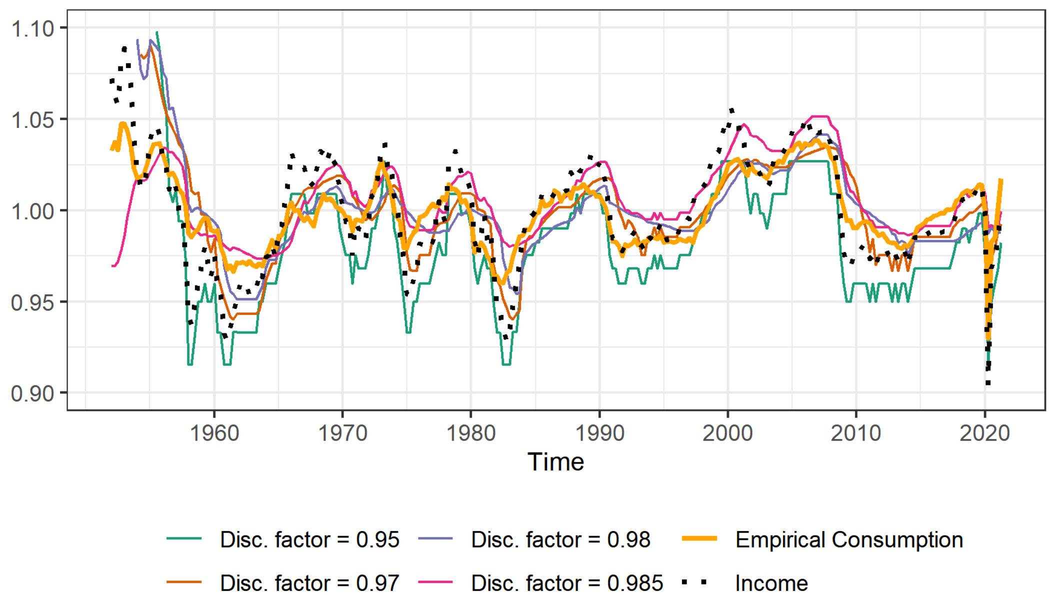

(b) Assets

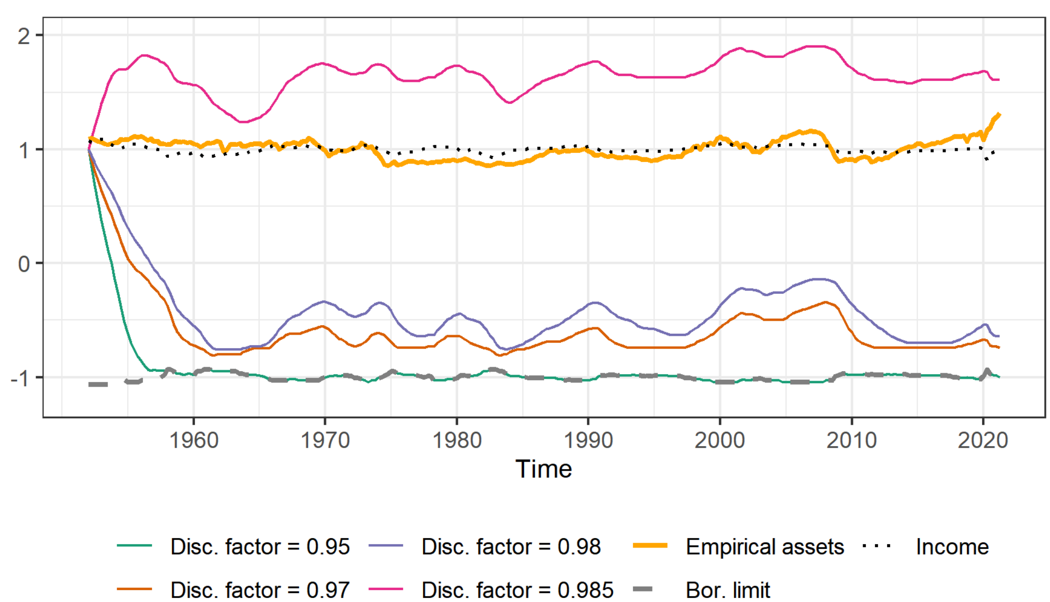

NOTE: Income levels are based on an AR1 simulation of the log of the de-trended and normalized empirical income series using OLS estimates. We then use 12 income levels equal to the 1st, 5th, 10th, 15th, ..., 95th, 99th percentile of the similated series to obtain a discretized version of the original de-trended and normalized income series and estimate a corresponding transition matrix based on the simulated series. We thus obtain a discrete Markov process with 12 income levels. The dotted income profile shows to the original non-discretized income values while the consumption profiles are based on the discretized Markov process.

this analysis is that consumption can follow income very closely even if decision making is forward looking when the agent values future opportunities based on a rough approximation. The asset profiles in Panel (b) initially move towards a stationary level and then exhibit a very low fluctuation. Although we show results for a wide range β values, the figure shows that asset levels vary only moderately with β, in particular in comparison to the RE case discussed below.

Panel A in Table 1 shows correlation coefficients for simulated and empirical consumption and asset series for various values of β. We show these correlations for various models of expectation formation as indicated in the column header. The first two columns refer to similarity-based forecasts, the third and forth column show results for constant forecasts equal to the mean of income. The next two columns show results for a forecast based on naive extrapolation, i.e. $y^*_{t+1|t} = y_t$. In the last two columns, expectations come from the University of Michigan Survey of Consumers. Specifically, we use the response to the survey question about expected change in the financial situation in a year. The possible answers to this question are "better off", "same", "worse off". The university of Michigan provides a score equal to the percentage of those answering "better off" minus those answering "worse off" plus 100. We rescale this measure to the range of our empirical income series and use this an empirical measure of $y^*_{t+1|t}$. The correlation coefficients are calculated from 1965 onward and exclude the initial adjustment phases of the model-based series. Correlations are highest for historical analogies and simple extrapolations. For most β values, the correlations with empirical consumption range around 0.85 for both similarity-based and simple extrapolation forecasts. For assets, the correlations range slightly above 0.5. Correlations are somewhat lower for constant forecasts, and significantly lower for the Michigan survey expectations in spite of the fact that these reflect real-time empirical expectations. The fact that for a range of β values there is no great difference between constant forecast models and historical analogies or naive extrapolations suggests that expectations are not the main drivers of FU decisions but only secondary to the FU value function. However, they are clearly not irrelevant, as the Michigan results show.

The fact that empirical consumption closely moves together with income invites comparison to a rule of thumb model in which the agent consumes a fixed proportion of income. However, such a model would tend to show a low correlation with empirical assets. To see this, assume that the agent chooses $c_t = y_t$ for all t. Then assets would be virtually determined by their initial levels and would increase in proportion with returns (for positive asset levels), which

Table 1: Correlations between simulated and empirical series for FU and RE models

Panel A: FU model

	Historical analogies		*Constant forecast*		*Simple extrapolation*		*UoM Survey*	
Beta	Cons	Assets	Cons	Assets	Cons	Assets	Cons	Assets
0.5	0.842	-0.383	0.862	0.536	0.860	-0.535	0.770	-0.334
1.0	0.856	0.047	0.848	0.535	0.872	-0.532	0.792	-0.15
1.5	0.859	0.478	0.834	0.533	0.873	0.534	0.800	0.011
2.0	0.857	0.527	0.821	0.530	0.869	0.529	0.794	0.124
2.5	0.852	0.530	0.809	0.523	0.864	0.524	0.789	0.198
3.0	0.846	0.525	0.798	0.516	0.857	0.516	0.783	0.247
3.5	0.840	0.516	0.788	0.507	0.850	0.507	0.776	0.272
4.0	0.833	0.505	0.777	0.496	0.843	0.496	0.769	0.288
6.0	0.804	0.451	0.741	0.447	0.814	0.447	0.719	0.261
8.0	0.768	0.379	0.707	0.393	0.784	0.393	0.594	0.147
10.0	0.707	0.284	0.670	0.338	0.750	0.336	0.432	0.034

Panel B: RE model

Disc. factor	Consumption	Assets
0.910	0.862	-0.538
0.930	0.861	-0.538
0.950	0.862	-0.538
0.970	0.757	0.396
0.980	0.694	0.251
0.985	0.785	0.311

NOTE: The table shows correlation coefficients for empirical and model-based consumption and asset values for various β levels for the FU model (Panel A), and various discount factor levels for the RE model (Panel B). The correlation coefficients are calculated from 1965 onward and exclude the initial adjustment phases of the model-based series. UoM refers to the University of Michigan Survey of Consumers.

are assumed constant in the model. This would lead to a very low correlation with empirical assets. By contrast, for the FU model, even in the case of $\beta = 2.2$ where assets approximately stay at the initial level (see Panel (b) of Figure 6), the correlation with empirical assets is around 0.45 (not shown in Table 1).

Figure 7 shows predicted consumption and asset profiles for the RE model for various levels of discount factors. The underlying income process is derived from the empirical profile shown in Figure 1 as follows. We first estimate an AR1 process for the log of the income series in Figure 1 using OLS.[17] We then simulate a time series of 100,000 periods according to an AR1 process with parameters equal to the OLS estimations and assuming normal errors. The resulting series is converted to levels. We then select 12 discrete levels equal to the 1st, 5th,

[17] Higher-order AR estimations do not increase model fit in any significant way.

10th, 15th, ..., 95th, 99th percentile of the simulated income series (in levels) and discretize the series from Figure 1 using these levels. We finally estimate a transition matrix based on the simulated series. In the RE model, the agent knows the resulting Markov process and considers it as the true data-generating process. The discrete approximation of income is visible in the consumption profile for a discount factor of 0.95 in Figure 7, where consumption moves parallel to income. The income profile depicted in the figure by the black dotted line is the original continuous version.

The RE profiles exhibit patterns that are very similar to those in Figure 5. Again, we set initial assets equal to 1. For discount factors below 0.95, assets hit the borrowing constraint very soon and then constantly stay there. As a result, consumption moves fully in-tandem with income. In Panel (a), consumption and income do not overlap, however, because the agent has to deduct interest payments before consuming the rest of income. When the discount factor increases to 0.98, consumption profiles gradually become smoother. Since assets remain rather close to the borrowing limit, there is still substantial comovement between consumption and income. Parallel to the situation for three income levels, predictions become highly sensitive to small changes in the discount factor when the latter exceeds 0.98. While assets fluctuate around -0.5 for a discount factor of 0.98, they fluctuate around 1.7 for a discount factor of 0.985. For still higher discount factors, assets quickly increase to very high levels.

Panel B in Table 1 shows the correlation between predicted consumption and asset values and the respective empirical values for the RE model. For low discount factors, the correlations with consumption are marginally higher than for the FU model. However, the correlation with the empirical asset series is negative for those discount factors. For higher discount factors, when the correlations with empirical assets turn positive, the correlations of predicted and empirical consumption drop significantly below the corresponding values for the FU model. The reason is that the RE model predicts already too much smoothing for those discount rates. Overall, the FU model captures the empirical pattern better than the RE model. This is not due to an a priori higher degree of flexibility that would make it easier for the FU model to match any empirical pattern. The degree of flexibility of both the FU and RE models is similar. In fact, the constant forecast version of the FU model is particularly parsimonious and captures the patterns in the data almost as well as the similarity-based version. The main driver of the increased empirical fit of the FU model is the FU value function and the implied rebalancing of marginal values of current consumption and future opportunities. It is this feature that also generates a high degree of comovement of FU consumption series with income. However, in spite of the favorable numbers in Table 1 it should be mentioned that the FU model in its current form somewhat underestimates empirical consumption smoothing. Finding parsimonious modifications of the FU model that improve its predictions with respect to consumption smoothing is an important topic for future research.

7 Conclusion

The starting point of this paper has been the observation that, for real-world actors, the data generating process (DGP) of the economic environment may sometimes not be pinned down to a narrow range. Moreover, actors may anticipate that it may change in unpredictable ways. We refer to such an environment as one of fundamental uncertainty (FU). In this paper, we present a model of how agents may make forward-looking decisions in an FU environment. The key concepts in this model are: an FU value function that is unspecific about DGPs; and forecasting of the argument of that value function by means of historical similarities. Technically, the latter is implemented using the knn framework from statistical learning. We compare this to a benchmark of a constant forecast. We consider the consumption and savings choice of a representative agent earning labor and transfer income and who is subject to borrowing constraints. Both empirical and simulated consumption profiles for the FU model show a high correlation with income. However, there is also a moderate degree of consumption smoothing. Overall, the FU model fits empirical data quite well. Unlike in a rule-of-thumb model, choices in the FU model are forward-looking and derived from behavioral goals. Ultimately the FU model is rooted in a standard logic of optimal intertemporal choice. It highlights how forward-looking agents may end up showing near-rule-of-thumb behavior due to approximating the value function in a way that is feasible under FU.

We compare our predictions to a simple rational expectations (RE) benchmark with logarithmic utility that shares a similar degree of flexibility. For relatively low discount factors, the RE model can also predict consumption profiles with a high correlation with income. This comes with a prediction of very low asset holdings and frequently binding borrowing constraints. At the same time, the correlations of these asset series with their empirical counterpart are negative. For higher levels of discount factors, the RE model predicts too much smoothing. Moreover, the predictions become extremely sensitive to the discount factor. Overall, the predictions of the FU model matches the empirical patterns better than the RE model.

The FU model is not perfect. It underpredicts smoothing in the aggregate consumption series. Exploring how the FU value function may be adapted in a parsimonious way to improve predictions regarding consumption smoothing is an important topic for future research. Further important research topics include an adaptation of the FU value function and expectation formation to a microeconomic environment; they also include exploring the implications for equilibrium models and the potential arising of multiple equilibria due to the expectation formation process.

References

Acemoglu, Daron (2021). "Harms of AI". National Bureau of Economic Research Working Paper 29247.

Acemoglu, Daron and Pascual Restrepo (2021). "Tasks, Automation, and the Rise in US Wage Inequality". National Bureau of Economic Research Working Paper 28920.

Autor, David, David Dorn, and Gordon H Hanson (2021). "On the Persistence of the China Shock". National Bureau of Economic Research Working Paper 29401.

Binswanger, Johannes (2011). "Dynamic Decision Making with Feasibility Goals: A Procedural-Rationality Approach". *Journal of Economic Behavior and Organization* 78, pp. 219–228.

— (2012). "Life Cycle Saving: Insights from the Perspective of Bounded Rationality". *European Economic Review* 56, pp. 605–623.

Binswanger, Johannes and Manuel Oechslin (2021). "The Economics of Beliefs under Fundamental Uncertainty".

Binswanger, Johannes and Martin Salm (2017). "Does Everyone Use Probabilities? The Role of Cognitive Skills". *European Economic Review* 56, pp. 73–85.

Carroll, Christopher D et al. (2020). "Sticky Expectations and Consumption Dynamics". *American Economic Journal: Macroeconomics* 12.3, pp. 40–76.

Cette, Gilbert, Lorraine Koehl, and Thomas Philippon (2019). "Labor Shares in Some Advanced Economies". National Bureau of Economic Research Working Paper 26136.

Gabaix, Xavier (2014). "A Sparsity-Based Model of Bounded Rationality". *The Quarterly Journal of Economics* 129.4, pp. 1661–1710.

— (2016). "Behavioral Macroeconomics via Sparse Dynamic Programming".

Gabaix, Xavier et al. (2006). "Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model". *American Economic Review* 96.4, pp. 1043–1068.

Hansen, Lars Peter and Thomas J Sargent (2008). *Robustness*. 2. ed. Princeton University Press.

Hansen, Lars Peter and Thomas J Sargent (2010). “Wanting Robustness in Macroeconomics”. In: *Handbook of Monetary Economics*. Vol. 3. Elsevier, pp. 1097–1157.

— (2022). “Structured Ambiguity and Model Misspecification”. *Journal of Economic Theory* 199. Symposium Issue on Ambiguity, Robustness, and Model Uncertainty, pp. 1–32.

Malmendier, Ulrike and Stefan Nagel (2016). “Learning from Inflation Experiences”. *The Quarterly Journal of Economics* 131.1, pp. 53–87.

Malmendier, Ulrike and Leslie Sheng Shen (2018). “Scarred Consumption”.

Segal, Michael (2018). “Automatic Pilots”. *Nature* 563.

Shiller, Robert J. (2017). “Narrative Economics”. *American Economic Review* 107.4, pp. 967–1004.

— (2019). *Narrative Economics. How Stories Go Viral and Drive Major Economic Events.* Princeton University Press.

Sutton, Richard S and Andrew G Barto (2018). *Reinforcement learning: An introduction.* MIT press.

The Economist (2022). “A Turning Point”. *The Economist* January 29, 2022.

Woodford, Michael (2019). “Monetary Policy Analysis when Planning Horizons are Finite”. *NBER Macroeconomics Annual* 33.1, pp. 1–50.