BULLETIN OF POLITICAL ECONOMY
15:1 (2021): 1-43

The Euler Equation Approach: Critical Implications of Recent Developments in the Theory of Intertemporal Choice

DARIA PIGNALOSA*

Abstract: The paper deals with the development of the mainstream theory of intertemporal choice based on the idea that individuals choose consumption and saving in order to maximize lifetime expected utility subject to an intertemporal budget constraint. This analytical framework rests on the theories developed in the 1950s by Modigliani and Friedman and revisited in the 1970s with the introduction of the rational expectations assumption. Starting with Hall (1978), the literature focused on testing the model relying on the first order conditions of the optimization problem faced by the consumer — the Euler equation — and a number of empirical puzzles arose. Therefore, the subsequent decades were dedicated to progressively modifying the original model so as to render it able to explain the data. The paper argues that the introduction of highly specific assumptions, needed to reconcile theory and empirical evidence, has affected the generality of the implications that can be drawn from the model. Furthermore, the ever more substantial departure from the original formulation seems to have resulted in a gradual abandonment of the very premises on which the neoclassical approach to consumption analysis was built.

Keywords: consumption function; Euler equation; intertemporal preferences; life-cycle model; household saving.

JEL classification: D01; D14; D15; D91; E21.

INTRODUCTION

The modern theory of intertemporal choice rests on the idea that individuals allocate resources over time so as to maximize lifetime expected utility subject to an intertemporal budget constraint. The origins of this approach are to be found in the Life Cycle-Permanent Income Hypothesis, which was put forward by Modigliani and Brumberg (1954) and Friedman (1957)

* Roma Tre University. Email: daria.pignalosa@uniroma3.it

during the debate on the consumption function that followed Keynes's (1936) contribution. In direct opposition to Keynes's analysis of aggregate consumption behaviour, Modigliani and Friedman developed an analytical framework based on individual optimization. In the 1970s, the Life Cycle-Permanent Income Hypothesis was revised with the introduction of the rational expectations assumption and Hall (1978) proposed to use the first order conditions of the intertemporal optimization problem faced by the consumer — the Euler equation — to derive a set of conditions suitable for empirical verification. Starting with Hall, the literature focused on testing the model of intertemporal utility maximization relying on the Euler equation and a new approach to the analysis of saving behaviour was established. Currently, the consumption Euler equation is not only the core of the theory of intertemporal choice but also a crucial component of dynamic general equilibrium models and hence one of the most popular tools of macroeconomic analysis.

The early empirical tests of the formulation proposed by Hall found several results that contradicted theoretical predictions and a number of empirical puzzles arose. The subsequent literature has attempted to provide an interpretation of these troublesome findings and to progressively modify and enrich the original formulation of the model so as to render it able to explain the data. This led to the more sophisticated versions of the intertemporal utility maximization model currently employed. In the literature on consumption, the link between empirically-oriented contributions and theoretical developments is indeed quite strong: the development of the approach has been in fact overwhelmingly influenced by the need to find versions of the model that could account for observed facts.

This paper provides a survey of the contributions that have highlighted the various empirical problems faced by the modern theory of intertemporal choice and of those that have proposed extensions and refinements of the basic model with the purpose of reconciling theoretical predictions and empirical evidence. We do not aim for a comprehensive reconstruction of such an extensive literature. Over the years, some of the most important scholars in the field have provided thorough reviews (Deaton 1992, Browning and Lusardi 1996, Browning and Crossley 2001, Attanasio and Weber 2010, Jappelli and Pistaferri 2017) to which the interested reader is referred. The aim of this paper is rather to present the evolution of the mainstream analysis of consumption over the last decades with a view to highlighting some critical aspects of the current stage of development of this field of research. In particular, we shall inquire whether, in the richest versions of the model currently employed, the necessary introduction of highly specific assumptions aimed at achieving the sought-for adherence to empirical evidence affects the generality of the implications that may be drawn from the model. Some problematic issues seem indeed open in this respect. Furthermore, we shall argue that the ever more substantial departure from the original formulation of the model has resulted in a progressive abandonment of the very premises on which the mainstream approach to consumption analysis was built, and that this process has occurred in the absence of the theoretical discussion that such a rupture would have required.

In the next section we shall present the formulation of the intertemporal utility maximization model adopted by Hall (1978). We shall then illustrate, in the subsequent three sections, the main empirical puzzles that arose when the model was taken to the data and the response to such troublesome evidence, involving changes to the original formulation. Some critical remarks on the recent development of the mainstream theory of consumption are presented in the final section. Throughout the paper, we shall not dwell on the technical aspects of the empirical tests performed and of the theoretical models proposed but rather seek to convey the gist of the various contributions discussed; indeed, a secondary purpose that distinguishes the present survey from those already available lies in the attempt to make the reconstruction readily accessible to readers who are unfamiliar with modern consumption literature or with the econometrics of intertemporal choice.

THE CONSUMPTION EULER EQUATION IN ITS ORIGINAL FORMULATION

As already mentioned, the roots of the modern approach to consumption theory lie in the life-cycle model proposed by Modigliani and Brumberg (1954) and in the permanent income hypothesis formulated by Friedman (1957). The main implication of both these theories is that individuals prefer smooth paths of consumption over the life cycle and thus use borrowing and saving to prevent income variability from causing large fluctuations in consumption. The idea is that consumers form estimates of their lifetime resources and set current consumption to the appropriate fraction of that estimate. Modigliani conceives the estimate in terms of wealth, so that the fraction is the annuity value of lifetime wealth, while Friedman conceives it in the form of permanent income, so that the fraction is very close to one. Moreover, the permanent income hypothesis envisages an individual facing an infinite-time horizon and is particularly focused on the consumer's attempts to smooth short-run fluctuations in income, whereas the life-cycle model considers a finite horizon and is therefore more oriented to the study of retirement saving. However, the literature often refers to the two theories as a single analytical framework, called Life Cycle-Permanent Income Hypothesis (LCPIH), because the key insight of the theories is the same: the resources that an individual allocates to consumption in any time period depend on his lifetime resources and not on the level of his current income; accordingly, saving reflects the discrepancy between current and permanent income, that is "transitory income". This leads to conclude that the traditional Keynesian demand management, through transitory tax policy or other transitory income boosting measures, is ineffective. And, clearly, the policy implications of the *General Theory* were, along with the lack of microeconomic foundations, the reason why the Keynesian analysis of aggregate consumption appeared hardly reconcilable with the neoclassical approach.

As a matter of fact, Modigliani's and Friedman's contributions nourished the debate on the consumption function that originated from the work of Keynes (1936). As is well known, Keynes's analysis of consumption and saving behaviour was to be translated in the so-called "Keynesian consumption function" according to which aggregate consumption depends on current aggregate income, with a decreasing average propensity to consume and a marginal propensity to consume lower than unity. When researchers tried to verify empirically the relationship between consumption and income, several findings appeared at odds with the Keynesian consumption function and, what is more, the evidence from different studies seemed scarcely reconcilable. In particular, it proved difficult to explain the coexistence of a constant average propensity to consume in time series data with the marginal propensity to consume being lower than the average propensity in cross-section analyses. The data appeared to suggest that consumption changes are less than proportional to cyclical income changes but at the same time proportional to income changes due to long-run growth. Many scholars entered the debate both with empirical contributions and with possible explanations of the observed behaviour. Compared to competing theories, the LCPIH was able to better explain the main stylized facts about consumption[1] and, perhaps just as importantly, it provided a theory with microeconomic foundations based on individual optimization and therefore consistent with the premises of the neoclassical approach. Thus, when the LCPIH was proposed, a general consensus was reached and it became the standard theory of consumption.[2]

In the 1970s, the assumption of rational expectations made its way into mainstream economic analysis and the theory of intertemporal utility maximization was modified accordingly. Hall (1978) examined the stochastic implications of the so-called Rational Expectations-Permanent Income Hypothesis (REPIH) and proposed a new econometric approach to test the predictions of the theory.

The framework considered by Hall is that of an individual[3] who maximizes lifetime expected utility, i.e.,

$$\mathbb{E}_t\left[\sum_{j=0}^{T-t} \beta^{t+j}\, u\left(C_{t+j}\right)\right],$$

where C_t is the level of consumption in period t; $u(\cdot)$ is the within-period utility function, which is continuous, increasing, and strictly concave; $\beta = 1/(1+\delta)$, is the discount factor, which is positive and smaller than unity because the rate of time preference, δ, is positive; $\mathbb{E}_t$ is the expectations operator that takes expectations of variables conditional on the information available at time t. The consumer has, in each period, a stochastic exogenous labour income Y_t, he can move resources over time, and is assumed to maximize the expected utility of consumption subject to an intertemporal budget constraint and a terminal condition on wealth:

$$\max_{C_{t+j}} \mathbb{E}_t\left[\sum_{j=0}^{T-t} \beta^{t+j}\, u\left(C_{t+j}\right)\right]$$
$$\text{s.t.}\quad A_{t+1} = (A_t + Y_t - C_t)(1+r)$$
$$A_T \geq 0$$

where A_t is wealth at time t and r is the rate of interest, assumed to be constant over time. The first order conditions of the optimization problem deliver the Euler equation for consumption:

$$u'(C_t) = \frac{1+r}{1+\delta}\mathbb{E}_t[u'(C_{t+1})] \qquad (1)$$

The Euler Equation can be seen as the usual optimality condition equating marginal costs and marginal benefits: the left-hand side of equation (1) represents the marginal utility cost of foregoing one unit of consumption at time t, while the right-hand side represents the marginal utility benefit from investing that unit at time t and consuming the proceeds at time t + 1.

The optimality condition of the intertemporal maximization problem states that the consumer aims at keeping the discounted expected marginal utility of consumption constant over time and thus implies that marginal utility behaves as a random walk (with trend): "no information available in period t apart from the level of consumption, C_t, helps predict future consumption, C_{t+1}, in the sense of affecting the expected value of marginal utility. In particular, income or wealth in periods t or earlier are irrelevant, once C_t is known" (Hall 1978: p. 974).

Hall's idea was to derive empirically testable restrictions from the consumption Euler equation in order to test the validity of the model. However, in its more general formulation, the Euler equation has no testable implications, since its predictions concern the marginal utility of consumption, which is not observable. Taking the model to the data requires to specify individual preferences, so as to derive predictions about consumption itself rather than its marginal utility. In order to perform an empirical test of the model, Hall adopted a quadratic within-period utility function, according to which

$$u(C_t) = -\frac{1}{2}(\bar{C} - C_t)^2$$

where $\bar{C}$ is the so-called "bliss level" of consumption. He therefore derived the following formulation for the Euler equation:

$$\mathbb{E}_t[C_{t+1}] = \alpha + \gamma C_t \qquad (2)$$

with $\alpha = \bar{C}(r - \delta)/(1 + r)$ and $\gamma = (1 + \delta)/(1 + r)$.

According to equation (2), the sign of the difference $r - \delta$ is all that is needed to establish whether consumption is increasing $(r > \delta)$ or decreasing $(r < \delta)$ over the life cycle. If we set $r = \delta$, thereby assuming that impatience is exactly offset by the incentive to postpone consumption, then equation (2) reduces to

$$\mathbb{E}_t[C_{t+1}] = C_t \qquad (3)$$

and we get the well-known random-walk proposition,

$$\Delta C_{t+1} = \varepsilon_{t+1}, \qquad (4)$$

where the expectational error, because of the assumption of rational expectations, is orthogonal to any information available to the consumer, i.e., $\mathbb{E}_t[\varepsilon_{t+1}] = 0$.

Equations (3) and (4) imply that, in each period, expected next-period consumption equals current consumption and the change in consumption between any two periods equals the change in the individual's expectations of his lifetime resources. In other words, the model of intertemporal choice put forward by Hall predicts that consumers will adjust their consumption plans only when they receive new information about their lifetime resources. As a consequence, an anticipated change in income should not affect consumption at the time it occurs, because the consumer would have already incorporated the expectation of the income change in his optimal consumption plan when the information first became known. As for changes in income that are not foreseen, consumption should react one-to-one to permanent shocks and should not be affected appreciably by transitory shocks.

In his contribution, Hall (1978) provided an empirical test of the implication that, conditional on current consumption, other current variables, including income, should not help in predicting future consumption, or, equivalently, no variables known in period $t-1$ and earlier should be correlated with changes in consumption between $t-1$ and t. He found that on U.S. macro data neither lagged consumption nor lagged income terms are significant, thus corroborating the main prediction of the model. However, he did find that lagged stock market prices have some explanatory power.[4] Shortly after Hall's contribution, the empirical research found several, more substantial, rejections of the random-walk proposition.

But before moving on to discuss this literature, we need to dwell on the assumptions on which the random-walk proposition rests. Although a thorough critical discussion is beyond the scope of this paper, it is important to be aware of the (often implicit) assumptions underlying the original formulation of the intertemporal optimization problem.[5] As we shall see, the subsequent developments of the theory, driven by the need to build versions of the model consistent with the empirical evidence, arise precisely from the modification of some of those assumptions, with the purpose of obtaining different predictions from the Euler equation. Furthermore, it is one of our purposes to argue that, while some of the original assumptions can be relaxed "easily", some others affect the nature of the model.

The framework that is implied in the Euler equation obtained by Hall is the following. Life is supposed to be finite and of known length T. The only source of uncertainty is that deriving from labour income, which is exogenous, so that the individual does not make labour supply choices. It is assumed that there is a single financial asset that the consumer can use to move resources across periods, which yields an interest rate that is known and constant. Credit markets are assumed to be perfect, i.e., lending and borrowing rates are the same and there is no credit rationing.

A number of hypotheses about preferences are implied even by the very general Euler equation (1), which does not require specifying a functional form for the utility function. The consumer is assumed to be rational and fully informed about prices and financial opportunities. He forms rational expectations about future incomes and hence about future marginal utility from consumption, which implies that the expectational error for the Euler equation is independent of the variables in the information set of the consumer. The consumer maximizes expected utility, which means that preferences satisfy the completeness, transitivity, continuity, and independence axioms and, as a consequence, utility is additively separable across states of nature. The within-period utility function is monotone increasing and strictly concave or, analogously, marginal utility is positive and decreasing. This ensures that the dynamic programming problem is solved by a value function that is itself continuous, concave, and strictly increasing, and provides a unique decision function. Preferences are also time-invariant, i.e., the within-period utility function is the same in each period of the time horizon. In addition, preferences are additively separable over time, which means that, in each period, marginal utility is independent from consumption in any other period. This precludes the consumer from deriving utility from goods over many periods and thus rules out habit-forming and durable (or storable) commodities. Consumers are ‘atomistic agents’, i.e., their preferences are independent from the choices of other consumers. A discount factor is applied to future utility flows, which implies that a given amount of consumption in the current period is valued more highly than the same amount of consumption next period, that is, the consumer is “impatient”. Plus, since the discount factor is constant, exponential discounting is assumed, so that the marginal rate of substitution between consumption at any pair of points in time depends only on how far apart those two points are. Utility is derived from per-period “consumption”: this can either mean that there is a single homogeneous consumption good or that utility is defined over a composite consumption good, in which case either constant relative prices or within-period homotheticity needs to be assumed. No variable apart from consumption appears as an argument of the within-period utility function, which implies either that the individual derives utility only from consumption or that preferences over consumption and any other variable affecting utility are additively separable. In particular, either the individual derives no utility from leisure or the marginal utility of consumption is not affected by the work status.

This set of assumptions leads to Euler equation (1) stating that the discounted marginal utility of consumption should be constant over time. In Hall’s contribution, as indicated above, it is further assumed that preferences are quadratic[6], which implies that marginal utility is linear. This leads Hall to conclude that not only the marginal utility of consumption evolves according to a random walk with trend, but, “to a reasonable approximation, consumption itself should evolve in the same way” (Hall 1978: p. 971).

For several reasons, the Euler equation is a very convenient empirical tool. First, it allows to take the model to the data without explicitly modelling all the details of the stochastic context in which the consumer operates. Second, the estimation of an Euler equation requires to have data on consumption and interest rates in different periods but does not require to observe wealth. Third, and most important, it allows to test the validity of the model (and, as we shall see, to estimate preference parameters) without having to solve explicitly the dynamic optimization problem of the consumer. That would require combining the Euler equation with the intertemporal budget constraint; however, a closed-form solution for current consumption as a function of the relevant state variables can be obtained only when strong assumptions about the nature of uncertainty are made and preferences are assumed to be either quadratic or exponential.[7] When the dynamic optimization problem cannot be solved analytically, it is necessary to use numerical solution techniques (such as in Deaton 1991 and Carroll 1992), but these require a complete specification of the economic framework faced by the consumer, which implies the need for additional and usually very restrictive assumptions. Through the Euler equation, instead, one can address the analysis of a problem that is analytically intractable by eluding the need of finding closed-form solutions.

It should be emphasized that the consideration of an equilibrium relationship such as the Euler equation, while imposing some restrictions on the dynamics of consumption, precludes the possibility of predicting the *level* of consumption or the way consumption moves in response to changes in the economic framework. However, while this can arguably be considered as a severe drawback, it is also what allows to overcome the problems highlighted by Lucas (1976), who pointed out that the relationship between consumption, income, and interest rates depends on the wider macroeconomic context and is therefore not stable over time. By focusing on the individual optimization problem, the Euler equation approach shifts the attention from the aggregate consumption function to the consumers' attempts to maximize their utility and is hence immune to the Lucas critique.

Finally, the result found by Hall is in line with the basic insight of both Modigliani's and Friedman's theories: that individuals *smooth consumption* across the life cycle. All these features contributed to make the Euler equation approach the standard approach to consumption theory.[8]

EARLY TESTS OF THE EULER EQUATION: THE EMERGENCE OF EMPIRICAL PUZZLES

In this section, we shall briefly illustrate the main puzzles arisen when the consumption Euler equation was taken to the data. Indeed, starting with the work of Hall (1978), the literature on consumption has focused primarily on the empirical tests of the intertemporal utility maximization model based on Euler equations estimation. The main predictions submitted to empirical validation concern the behaviour of consumption in response to income changes: according to Hall's model, consumption should not react to anticipated changes in income; on the contrary, it should react to income changes when these are not foreseen. In addition, the behaviour of consumption across the life cycle should be independent from that of income, i.e., the hump-shaped profile of income should not be reflected on consumption but on saving. Since the Euler equation allows to estimate the coefficients of the utility function, the empirical literature has also engaged in determining the values to be attributed to preference parameters. Finally, the model of intertemporal choice has also implications for asset prices, so that some of the empirical tests of consumption theory stem from the finance literature.

The response of consumption to income changes

The first empirical failure of the model which was encountered, which is also the one that has received the greatest attention, generating extensive empirical work over the last decades, is the "excess sensitivity" puzzle, initially highlighted by Flavin (1981). Unlike Hall, Flavin estimates jointly the consumption and income equations, thereby finding that aggregate consumption does in fact depend on anticipated income changes, contradicting the prediction that, if the variation of income is foreseen, it should already be incorporated in the behaviour of the consumer and not have any impact on consumption when it occurs. Flavin's test indicates that the U.S. consumption response to a variation of income is over three times the value predicted by the model, i.e., that consumption exhibits excess sensitivity to expected income changes.[9] During the 1980s, additional evidence hinting at an excess sensitivity of consumption was accumulated (see in particular the studies by Campbell 1987 and Campbell and Mankiw 1989). In fact, the attention that this puzzle has attracted comes as no surprise, since it suggests that consumption depends, at least partially, on current income, which is precisely what is implied by the Keynesian approach and what Friedman's notion of permanent income was intended to refute.

The evidence of excess sensitivity of consumption was associated with the symmetric but opposite empirical problem: not only consumption appeared to be too *sensitive* to predicted income changes, it also turned out to be too *insensitive* to unexpected changes in income. As we have seen, according to Hall's model, the consumer should not react to anticipated changes in income, but when the variation of income is not foreseen, this should affect consumption in the moment it occurs. Empirical evidence suggests instead that consumption does not react enough to unanticipated income changes, leading to excess smoothness.[10]

It is worth recalling that one of the main stylized facts that the permanent income hypothesis purported to explain is indeed the smoothness of consumption, i.e., the fact that changes in aggregate income are associated with relatively small changes in aggregate consumption. Friedman claimed that consumption is determined by permanent income, not by current income, and permanent income is smooth relative to current income. The evidence presented by Deaton (1986), West (1988), and Campbell and Deaton (1989) showed that permanent income is in fact less smooth than current income and that, as a consequence, the smoothness of consumption cannot be explained by the intertemporal utility maximization model.

Both the excess sensitivity and the excess smoothness puzzles have been substantially confirmed by the empirical work carried out in subsequent decades. Evidence of excess sensitivity has been found on micro as well as on aggregate data; it has also been observed that the reaction of consumption is usually stronger when expected income changes are large and positive, while small or negative changes appear to have relatively weak effects. As for the empirical tests of the response of consumption to unexpected income changes, these have found an important degree of smoothness in aggregate consumption; however, studies conducted on household-level data (starting with the work of Dynan 2000) have usually found no evidence of excess smoothness, thereby leading to the additional problem of reconciling the macro and the micro evidence.[11]

The relationship between income and consumption across the life cycle

As we have already said, the most important implication of the intertemporal utility maximization model is that the time path of income is irrelevant for consumption because individuals use borrowing and saving to smooth out income fluctuations. This applies not only to short-run (or high-frequency) fluctuations in income but also to longer-run (low-frequency) fluctuations related to changes in the work status.[12] Individuals should then borrow prior to labour market entry, accumulate wealth during the working life, and dissave in retirement. Against this prediction, the evidence suggests that life-cycle profiles of income and consumption track each other: they both increase during the first part of the life cycle, reach a peak at a roughly similar age, and decline afterwards. Such correlation has been highlighted by Thurow (1969) and confirmed subsequently by, among others, Browning, Deaton, and Irish (1985), but the most popular piece of evidence is the one provided by Carroll and Summers (1991), who show that for many countries and different educational and occupational groups of individuals, both income and consumption profiles are "hump shaped" and, additionally, the consumption profiles appear to be steeper for the groups of individuals that have steeper income profiles (such as those including more educated individuals).

Related to this empirical problem is the one referred to as the "retirement consumption puzzle", originally highlighted by Hamermesh (1984). Analysing the relationship between consumption and lifetime wealth, Hamermesh suggested that consumers retire with inadequate savings and are therefore unable to maintain a level of consumption consistent with the one afforded before retirement. Also, it has been observed that the elderly either continue to save in retirement or decumulate their wealth much more slowly than would be predicted by the intertemporal maximization model. Such "wealth decumulation puzzle" has been found, among others, by Mirer (1979), Diamond and Hausman (1984), and Bernheim (1987).

Clearly, for a theory whose basic insight is that consumers pursue smooth consumption paths, the evidence revealing a lack of smoothing across stages of life is a major failure. In addition to this, however, it has to be noted that the fact that consumption tracks income during the life cycle attests, just like the evidence of excess sensitivity, that current income exerts a larger influence on consumption than the theory would imply, thereby providing support for the Keynesian approach to explaining consumption behaviour.

The estimation of preference parameters

Within the intertemporal utility maximization model, the main characteristics of individual preferences are represented by a definite set of parameters. Such parameters are: i) the rate of time preference, which represents the tendency of individuals to attach more value to current than to future consumption (it measures the degree of impatience); ii) the coefficients of absolute and relative risk aversion, which represent the tendency of individuals to prefer a certain amount of consumption over an aleatory amount with the same expected value (they measure the price of risk); iii) the coefficients of absolute and relative prudence, which represent the tendency of individuals to save as a result of the presence of uncertainty (they measure the intensity of the precautionary motive for saving); and iv) the elasticity of intertemporal substitution (EIS), which represents the tendency of individuals to save in response to the incentive given by interest on saved sums (it measures the aversion to temporal fluctuations of consumption).

One of the advantages of the Euler equation approach that has generated more enthusiasm among scholars is that it allows to estimate the structural parameters of the utility function. However, in the forty years following Hall's contribution, Euler equation estimations have yielded inconclusive results, so that no consensus has yet been reached on the value of the parameters. As stressed by Carroll (2001: p. 1), "despite scores of careful empirical studies using household data, Euler equation estimation has not fulfilled its early promise to reliably uncover preference parameters".[13] This is rather troubling for, as we hope will appear clear in the following sections, without a complete specification of individual preferences the model of intertemporal utility maximization would prove devoid of definite empirical content, i.e., it would be unable to provide any indications about consumption and saving behaviour.

As a matter of fact, not only are the estimated values of the parameters quite variable, they also often appear rather implausible, to the point of generating in some cases authentic empirical puzzles. In the next sub-section we shall discuss the equity premium puzzle, related to the extremely high coefficient of relative risk aversion estimated in the consumption-based capital asset pricing model. Similarly implausible appear the small estimates of the parameter representing the precautionary motive for saving (see e.g. Guiso, Jappelli, and Terlizzese 1992 and Dynan 1993). Also, many contributions estimating the EIS suggest that consumption growth is completely insensitive to changes in interest rates (see in particular the influential works by Hall 1988 and Campbell and Mankiw 1989), which is quite unsettling for a theory according to which the consumption profile of individuals should be tailored so as to take advantage of the intertemporal price of consumption.

Furthermore, the estimates produced often appear in contradiction with the values adopted in other strands of literature. For example, in numerical simulations of models with precautionary saving (e.g. Zeldes 1989b and Deaton 1992) the assumed values of prudence are usually much larger than those obtained through Euler equations estimation. As pointed out by Eisenhauer (2000: p. 381), "evidently, either the simulations assumed unreasonable parameter values, or the empirical tests failed to capture the true extent of prudence. The widespread belief in prudent attitudes and precautionary motives tends to suggest the latter". Also, there is a notable discrepancy between the often close-to-zero values of the EIS obtained through Euler equations estimation and the dynamic macroeconomic literature, since general equilibrium models usually require an EIS close to one or even greater in order to match the evidence.[14]

Intertemporal choice and asset pricing theory

The intertemporal utility maximization model has direct implications for asset prices. These are drawn in the consumption-based capital asset pricing model (CCAPM) which links consumption and saving decisions to the choice of portfolio allocation. The CCAPM, derived in the works of Rubinstein (1976), Lucas (1978), and Breeden (1979), considers the optimization problem of an individual choosing his portfolio so as to maximize expected utility. The first order conditions for the optimal holding of assets relate consumption growth and asset returns and lead to an equation for the difference in the expected return on any two assets. Intuitively, this difference depends on the riskiness of the two assets and on the price of risk, with the latter being represented by the coefficient of relative risk aversion of a representative consumer.

The early empirical literature on consumption-based asset pricing reports several rejections of the implications of the model (Hansen and Singleton 1983, Grossman, Melino, and Shiller 1987, Breeden, Gibbons, and Litzenberger 1989). One of the most prominent empirical failures concerns the equity premium, i.e., the difference between the average return on the stock market and what is considered the risk-free interest rate, that is the return on short-term government bonds. Given the riskiness of equities, measured by the covariance of the excess stock return with consumption growth, the equity premium should be explained by the coefficient of relative risk aversion. But, since investors have historically received a very large premium for holding equities and at the same time equities appear to imply little risk, asset markets equilibrium requires consumers to have an implausibly high aversion to risk. This problem, known as the equity premium puzzle, has been first highlighted by Mehra and Prescott (1985), who employ data for the U.S. from 1889 to 1978 indicating that the return on equity has averaged 6.2% more than the return on short-term government bonds. Calibrating an asset pricing model with what they consider reasonable values of the preference parameters (including a coefficient of relative risk aversion below or equal to 10) they are not able to produce more than a 3.5% equity premium: the high equity premium observed would require consumers to have a coefficient of relative risk aversion in the 20-30 range.[15] The problem is not only that the degree of risk aversion consistent with the observed equity premium is too large to be believable. An additional puzzle — the so-called risk-free rate puzzle (Weil 1989) — arises because, as we shall see below, in the standard CCAPM with isoelastic preferences, the EIS is by definition the inverse of the coefficient of relative risk aversion. The link between the two parameters implies that consumers who are extremely averse to risk necessarily are extremely averse to intertemporal substitution as well. Now, just like individual attitudes towards risk are supposed to explain the equity premium, individual attitudes towards intertemporal substitution should explain the level of the risk-free rate. But the riskless rate of return consistent with a very low EIS is well in excess of the one that has been historically observed. In other words, even if one were willing to accept a patently implausible value for the coefficient of the utility function, it would still be impossible to replicate both the level of the risk-free rate and the equity premium within a CCAPM: an extremely large coefficient of risk aversion would explain the high equity premium but would be inconsistent with the low risk-free rate and vice versa.

REACTIONS TO THE EMPIRICAL PUZZLES: WRONG THEORY OR WRONG FORMULATION OF THE THEORY?

Overall, the empirical evidence accumulated in the first decade of the Euler equation approach appears to be essentially inconsistent with the basic implications of the model of intertemporal utility maximization. The reactions of scholars to these troublesome findings have been manifold and of different nature. We shall briefly present the main responses to the empirical puzzles in the present and the following section and devote the concluding section to a general discussion of the recent evolution of consumption research.

Utility maximizers vs rule-of-thumb consumers

Among the early responses to the unsatisfactory empirical performance of the intertemporal maximization model, the one provided by Campbell and Mankiw (1989) is particularly remarkable. In order to explain both the excess sensitivity and the excess smoothness puzzles, they suggest a model with two types of agents: the first type is a rational utility-maximizing consumer, whereas the second type simply consumes all disposable income and refrains from either borrowing or saving. In other words, Campbell and Mankiw (1989) suppose that a share of the population, rather than behaving according to the theory of intertemporal choice, follows a rule of thumb which says "consume current income". They estimate that a fraction as high as 50% of U.S. income goes to rule-of-thumb consumers and conclude that, compared to the intertemporal maximization model, "the old-fashioned Keynesian consumption function may […] provide a better benchmark for analyzing fiscal policy" (p. 210).[16] Such striking results are reinforced by a subsequent work in which, replicating the analysis for several other countries, Campbell and Mankiw (1991) find estimates of the fraction of income accruing to rule-of-thumb consumers ranging from 20% in Canada to nearly 100% in France. In the following years, further tests found results consistent with those of Campbell and Mankiw and it soon became common practice to refer to a share of "Keynesian" or "hand-to-mouth"[17] or "non-Ricardian"[18] consumers. For example, Weil (1992) shows that allowing for the existence of hand-to-mouth consumers contributes to the resolution of the equity premium puzzle and of the risk-free rate puzzle; Fuhrer (2000) and Kiley (2010) both consider frameworks that allow for alternative models and conclude that, in order to account for the observed excess smoothness of consumption, an important fraction of rule-of-thumb consumers must be incorporated; Bernheim, Skinner, and Weinberg (2001) examine wealth accumulation and consumption patterns before and after retirement and find that the data are hardly reconcilable with the intertemporal utility maximization model and more easily explained by rule-of-thumb behaviour.

Some papers seem to attempt to account for the existence of rule-of-thumb behaviour. For example, Lettau and Uhlig (1999) propose a model involving the possibility for an agent to "learn" to use a rule of thumb à la Campbell and Mankiw (1989) rather than a rule that prescribes to behave according to the model of intertemporal utility maximization, while Winter, Schlafmann, and Rodepeter (2012) perform numerical simulations of life-cycle saving decisions to calculate the utility losses associated with rule-of-thumb behaviour, suggesting that these can be relatively low.

As we shall see, rather than accepting such a patently unsatisfactory solution as the *ad hoc* introduction of a significant proportion of consumers who simply do not behave as prescribed by the intertemporal maximization model, most scholars in the field of consumption research have attempted to find alternative explanations for the empirical evidence. It must be noted, however, that models with rule-of-thumb consumers have never been abandoned and, what is more, the recognition of non-optimizing behaviour is not confined to the consumption literature. In an influential paper, Mankiw (2000) advocates the systematic incorporation of rule-of-thumb consumers in macroeconomic models so as to fully understand the policy implications of their presence. Following this suggestion, Galí, López-Salido, and Vallés (2004) extend the standard new Keynesian model to allow for the coexistence in the economy of an optimizing and a non-optimizing class of agents, thereby initiating a still lively tradition.

The original formulation as a special case

When the empirical failures of the intertemporal maximization model began to accumulate, the research focused on providing new interpretations of the available empirical evidence, improving the quality of the tests, and, at the same time, extending the baseline model in order to reconcile it with the data.

A first important step toward the rehabilitation of the intertemporal utility maximization model consisted in realizing that the latter need not be identified with the random-walk proposition. Early tests based on Euler equation estimation relied on the specific version of the model proposed by Hall (1978), which includes the key hypothesis of quadratic preferences. It is this particular parametrization of the utility function which (associated with the hypothesis of equality between the rate of time preference and the interest rate) leads to the random-walk proposition, i.e., which implies that not only the marginal utility of consumption but consumption itself should be constant over the life cycle. We shall come back to the implications of quadratic preferences when discussing the precautionary motive for saving; at this stage, what we need to stress is that some of the rejections originally found in the empirical literature can be regarded as failures of the specific version of the intertemporal utility maximization model proposed by Hall and not as violating the basic insights of the theory. The core of the theory is that individuals maximize lifetime expected utility subject to an intertemporal budget constraint, which implies that consumers should smooth marginal utility, not consumption. As a consequence, the empirical evidence suggesting that individuals have fluctuating paths of consumption does not necessarily contradict the model. Accordingly, researchers started taking account of variables other than consumption that are likely to affect marginal utility.

Attanasio and Browning (1995) suggest that the excess sensitivity of consumption to income may disappear when controlling for demographic variables (such as family size and number of children), whereas Blundell, Browning, and Meghir (1994) and Attanasio *et al*. (1999) show that allowing demographics to affect household preferences can generate hump-shaped consumption profiles of the kind found by Carroll and Summers (1991). A further issue connected with the preference specification relates to the impact of labour supply on marginal utility. In this respect, Attanasio and Weber (1995) suggest that if utility is a function of consumption and leisure and the two arguments of the function are nonseparable, then the saving behaviour is affected by anticipated changes in labour supply, so that failing to control for labour supply indicators may lead to spurious evidence of excess sensitivity. The evidence on the retirement consumption puzzle has been somewhat reconsidered as well, pointing out that some part of the drop in consumption at retirement may be planned and related to changes in the work status. In particular, home production and the cessation of work-related expenses may lead to a decline in consumption that does not imply an increase in marginal utility (Aguiar and Hurst 2005).

A further step in the attempt to rehabilitate the theory of intertemporal choice concerned the quality of the data used in estimating Euler equations. In particular, most of the early empirical tests of the Euler equation were performed using macroeconomic data, by regressing the aggregate rate of growth of consumption on the rate of growth of income. In the 1990s, the relevance of aggregation bias was explicitly addressed[19] and, also because of the availability of better consumption surveys, the use of aggregate data was abandoned. Concurrently, there have been significant advancements in the empirical strategies adopted to test the implications of the theory.[20]

Overall, improvements in the understanding of the implications of the theoretical model, in the interpretation of the empirical evidence, and in the quality of the empirical tests performed led to a considerable reshaping of the evidence contradicting the intertemporal utility maximization model. Nevertheless, the empirical validation of the theory was still disappointing, as witnessed by the assessment of Attanasio (1995: p. 40):

> While from a theoretical point of view the life-cycle model constitutes the most appealing and flexible framework with which to study consumption, there is no widespread agreement among economists on the empirical relevance of the model.

As a matter of fact, the main rejections of the theory, although reconsidered, survived this more mature stage of the empirical research and scholars kept feeling the need to look for extensions of the original model capable of reconciling the theory of intertemporal choice with the available data.

REACTIONS TO THE EMPIRICAL PUZZLES: EXTENSIONS OF THE MODEL

In the present section, we shall briefly illustrate the main departures from the original formulation of the intertemporal maximization model that have been proposed in the last decades. In some cases, it seems appropriate to regard the extensions of the model that have been proposed as refinements of the original formulation; however, we shall argue that, in several other cases, the extensions appear to clash more or less sharply with some of the analytical premises of the theory of intertemporal choice. For this reason, rather than presenting them according to the chronological order in which they were put forward, we shall group the various versions of the model on the basis of what seems to us a reasonable and hopefully useful classification. We shall thus begin with the model that involves liquidity constraints, which generates similar results to the one comprising a fraction of rule-of-thumb consumers but does not require to postulate non-optimizing behaviour. In this case it is in fact the framework in which the consumer operates that prevents him from satisfying the optimal conditions of the dynamic programming problem. We shall then move on to a set of models (the precautionary saving, the buffer-stock and the wealthy hand-to-mouth models) in which, as we shall see, it is the value assigned to preference parameters that seems to play a crucial role in explaining the observed consumption behaviour. Finally, we shall present those models that appear to mark a rather fundamental departure from the premises of the approach, in that they abandon some of the basic assumptions of the theory of intertemporal choice.

Liquidity constraints

Some of the early contributions attempting to provide an explanation for the evidence of excess sensitivity (Hayashi 1985, Zeldes 1989a, Deaton 1991) proposed to remove the hypothesis of perfect credit markets, introducing a limit to borrowing possibilities.

If the individual is willing to shift resources to the present so as to increase current consumption but he is not allowed to borrow, then the Euler equation will not hold as an equality, i.e., current marginal utility will be higher than discounted future expected marginal utility. In the periods in which the constraint is binding, the individual will consume his income, acting as a rule-of-thumb consumer of the kind pointed out by Campbell and Mankiw (1989). However, unlike allowing for rule-of-thumb behaviour, introducing credit market imperfections is entirely consistent with a model of intertemporal utility maximization, which is why the assumption that some consumers are liquidity-constrained has proven very convenient in the attempt at reconciling the theory with the empirical evidence. It should be noted, though, that liquidity constraints may help explain excess sensitivity to anticipated income *increases* but cannot explain why consumption reacts to anticipated income *declines*: when current income is low relative to permanent income and the consumer cannot borrow, consumption is set equal to current income, but when income is expected to decrease the individual can still save. As a matter of fact, the evidence suggests that consumption is much less responsive to expected income declines and this is usually interpreted as supporting the model with liquidity constraints. But, as far as direct tests of the hypothesis are concerned, "detecting liquidity constraints from the Euler equation is difficult in the extreme" (Jappelli and Pistaferri 2017: p. 93) since the actual binding of the borrowing limit is generally unobservable. In some contributions, the sample is split in order to examine separately the behaviour of households who are likely to be constrained (for example those with low liquid wealth or low income) and of those who are not (see e.g., Zeldes 1989a and Johnson, Parker, and Souleles 2006). Usually, these tests find that potentially constrained consumers react strongly to expected income increases but at the same time a significant consumption response also occurs among likely unconstrained households.

Precautionary saving

As we have already pointed out, when the utility function is quadratic, marginal utility is linear, so the expected marginal utility of consumption is the same as the marginal utility of expected consumption and hence an increase in uncertainty does not affect saving.[21] This case is known in the literature as the certainty-equivalence model, because it implies that the individual consumes the amount he would consume if his future incomes were certain to equal their expected values. If, instead, preferences exhibit prudence, i.e., the marginal utility of consumption is convex, then the individual engages in precautionary saving: prudence leads consumers to treat future uncertain income cautiously and not to consume as much currently as they would in the absence of uncertainty.[22]

In the late 1980s, the discontent with the certainty-equivalence model was already growing (Blanchard and Mankiw 1988) and, shortly after a precautionary motive for saving had been introduced in the Euler equation literature (Zeldes 1989b and Caballero 1990), the certainty-equivalence version of the model was definitively discarded. The quadratic parametrization for the utility function therefore disappeared from mainstream literature and was replaced by the assumption of exponential preferences (Caballero 1990) or by the now-standard isoelastic specification (Zeldes 1989b).

If individual preferences exhibit prudence, then consumption does not behave as a random walk, so the presence of a precautionary motive for saving has been advocated as an explanation for the evidence of excess sensitivity and for the observed relationship between the shape of consumption profiles and that of income profiles. In particular, precautionary saving can be relevant for consumers who are impatient and are therefore inclined to anticipate consumption to the current period: prudence may lead these consumers to postpone consumption in the face of uncertainty. Also, the precautionary motive for saving is likely to be important for individuals in the early part of the life cycle: a young consumer who expects his income to rise but displays prudence may be willing to borrow less than would be the case in a certainty framework or under quadratic preferences. This helps explain why consumption and income usually rise together in the early part of the life cycle. At the same time, if retirement savings become important at some point along the life cycle, they can be used to smooth out income fluctuations thereby making the precautionary motive less relevant. Prudence also interacts with liquidity constraints, because the possibility of future binding borrowing constraints provides an additional motive for saving.[23]

The model with precautionary saving is much less at odds with the empirical evidence than the original model, but is still unable to satisfactorily explain the observed consumption behaviour unless prudence is associated with other modifications to the original formulation (Jappelli and Pistaferri 2010).

Buffer-stock saving

When borrowing constraints are associated with an important precautionary motive for saving and a sufficient level of impatience, one obtains the so-called "buffer-stock model" proposed by Deaton (1991) and developed by Carroll (1992, 1997 and 2011). In such a model consumers balance the urge to consume prompted by impatience against the desire to save induced by prudence. As a result, buffer-stock consumers have a target level of liquid assets, above which impatience dominates and assets are run down, and below which the precautionary motive dominates and assets are accumulated.

The idea of savings as a "buffer stock" for contingencies can help explain the evidence of consumption tracking income across the life cycle:

consumers with the appropriate combination of liquidity constraints, impatience, and prudence will maintain a small amount of assets to protect consumption against transitory income changes but will not engage in long-horizon borrowing or saving. In other words, the buffer-stock model has the appeal of predicting consumption smoothing of short-run (high-frequency) income fluctuations while not implying (low-frequency) consumption smoothing across the life cycle.

Wealthy hand-to-mouth consumers

Kaplan and Violante (2014) and Kaplan, Violante, and Weidner (2014) propose a model featuring "wealthy hand-to-mouth" consumers, i.e., consumers who hold little or no liquid wealth despite owning sizable amounts of illiquid assets. These consumers, because of their portfolio configuration, respond to income changes as if they were liquidity constrained.

A distinct feature of the wealthy hand-to-mouth model is the coexistence of two assets: a low-return liquid asset (such as bank accounts) and a high-return illiquid asset (such as equities). By saving in the illiquid asset, consumers achieve a higher level of average lifetime consumption but, at the same time, they forego to smooth income fluctuations, which instead is allowed when wealth is held in the liquid asset. Through a suitable configuration of parameters, the model generates a remarkable fraction of consumers for whom it is optimal to hold their wealth in the illiquid asset and consume all of their income every period. The presence of such wealthy hand-to-mouth consumers, combined with those who consume hand-to-mouth because they have low incomes (and are subject to borrowing constraints), implies that a large fraction of the population is highly sensitive to income changes. This allows to produce empirical results similar to those obtained in a model with rule-of-thumb consumers without allowing for a fraction of the population to depart from the behaviour prescribed by an intertemporal utility maximization model.

Nonexpected-utility preferences

In the intertemporal utility maximization model, the elasticity of intertemporal substitution is constrained to be the inverse of the coefficient of relative risk aversion. This is due to the double additivity induced by the simultaneous assumptions of intertemporal separability (i.e., additivity over time periods) and expected utility (i.e., additivity over states of nature), which implies that the curvature of the utility function establishes the individual's attitude towards both time and state fluctuations in consumption.[24] This prevents a moderate aversion to risk and a strong aversion to intertemporal substitution to coexist, which is what would be required to replicate the empirical evidence.

In order to disentangle the two preference parameters, one of the sources of additivity must be dropped. Accordingly, Epstein and Zin (1989) and Weil (1989) developed a nonexpected-utility model of choice drawing on the work of Kreps and Porteus (1978). The resulting functional form, often called Epstein-Zin-Weil (EZW) recursive utility, has been increasingly employed in recent years. In particular, because of the flexibility which it provides, EZW utility has been adopted in the empirical work aiming at estimating preference parameters (starting with the contributions of Attanasio and Weber 1989 and Epstein and Zin 1991). It also found widespread application in the finance literature, in the attempt to solve asset pricing puzzles (see, e.g., Weil 1989 and Bansal and Yaron 2004).

It is important to stress that the assumption of EZW preferences implies the abandonment of the hypothesis of expected utility underpinning the mainstream theory of choice under risk. Indeed, EZW preferences are based on an axiomatization of behaviour which is alternative to expected utility and which, unlike the latter, involves a nonlinear aggregator of present and future utility (utility is recursively defined over current consumption and a certainty equivalent of future random utility).[25]

Habit formation

An extremely popular modification of the original model involves the removal of the hypothesis of intertemporal additivity of preferences, according to which marginal utility is, in each period, independent from consumption in any other period. This allows to consider what the Euler equation literature refers to as "habit formation", that is the dependence of current marginal utility on past consumption.

The model with habit formation has attracted much attention in the finance literature since it may provide a partial solution to the equity premium puzzle (Abel 1990, Constantinides 1990, and Campbell and Cochrane 1999). Habits increase the disutility associated with large declines in consumption, thereby inducing consumers to require a larger premium to hold risky assets relative to consumers with time-additive preferences. Habit formation can therefore explain the high equity premium observed without resorting to implausibly high levels of risk aversion. In addition, Meghir and Weber (1996) argue that when the intertemporal maximization model includes habit-forming preferences, the data show no evidence of excess sensitivity of consumption to expected income changes. However, the popularity of habit-formation models is primarily owed to their ability to explain the observed excess smoothness of consumption (Deaton 1986, Fuhrer 2000, and Kiley 2010). Indeed, the habit-formation specification of preferences provides consumers with a motive to smooth not only the level but also the change of consumption and thus leads them to respond gradually to unexpected income changes. Over the years, the assumption of habit formation has become so standard that the coefficient measuring the serial correlation in consumption growth, which in fact represents the excess smoothness of consumption, is usually understood as the habit-formation coefficient. In other words, evidence of excess smoothness is automatically interpreted as evidence of habit formation. It has to be noted, however, that the insensitivity of consumption to unexpected income changes is usually not observed in studies conducted on household-level data (e.g., Dynan 2000): excess smoothness is only found in aggregate consumption. On the contrary, habit formation implies consumption to be smooth at both the macro and the micro level.

Social preferences

Some contributions have extended the intertemporal utility maximization model in order to allow for social influences on consumption behaviour. The idea, which also underlies some of the models with "external" habit formation, dates back at least to the work of Veblen (1899) and had already been introduced in the theory of intertemporal choice by Duesenberry (1949) shortly before the LCPIH made its appearance.

The "keeping up with the Joneses" model prescribes utility to depend on the current average consumption of the individual's peers. This leads to a different saving profile from the one implied by the standard atomistic-agents model: if peers' consumption is on a growing trend then undersaving (or overborrowing) will occur, while in the opposite case the model gives rise to underconsumption (De Giorgi, Frederiksen, and Pistaferri 2020). Alternatively, utility can be assumed to depend on *past* average peer consumption (Ljungqvist and Uhlig 2000), in which case the model is supposed to capture "catching up with the Joneses" behaviour.

From imperfect information to rational inattention

Models of imperfect information processing have been explored as possible explanations for the evidence of excess smoothness. Pischke (1995) proposes a model in which consumers react optimally to their own income process but have incomplete or no information on economy-wide variables.

The idea is that gathering information is somehow costly for consumers (Lucas 1973) and, since uncertainty about the macroeconomic environment makes a negligible contribution to total individual uncertainty, households have little incentive to distinguish aggregate from idiosyncratic shocks. According to Pischke (1995: p. 807), "agents may simply not care enough about aggregate information because ignoring it is not very costly for (most) households." This implies that, when an aggregate shock occurs, consumers adjust their consumption only partially because they confuse it with a (less persistent) individual shock. Although such a model could in principle explain any amount of excess smoothness, when it is actually calibrated to fit the empirical evidence, it generates far less excess smoothness than exhibited in the data.

Ludvigson and Michaelides (2001) introduce Pischke's assumption of incomplete information on the source of income shocks into a buffer-stock model. They are able to generate excess smoothness, but in order to meet the degree of smoothness found in the data, they need to simultaneously admit an implausibly high degree of excess sensitivity.

More recently, Carroll *et al.* (2020) have managed to better replicate the empirical evidence by assuming that consumers have accurate knowledge of their personal circumstances but "sticky expectations" about the macroeconomy. The idea is not so much that consumers have imperfect information but rather that they are unable to attend to all information available and therefore display "rational inattention" (Sims 2003): their ability to allocate attention is limited and they allocate it optimally. More specifically, Carroll *et al.* (2020: p. 43) assume that "consumers […] only occasionally observe aggregate data." The lag in perception generates a lag in the response of aggregate consumption to aggregate developments. Unlike the model with habit-forming preferences, which implies consumption to have the same autocorrelation structure in micro and in aggregate data, the model with sticky expectations is consistent with both household-level and macro evidence, because the excess smoothness that it generates in aggregate consumption stems from aggregation of the microeconomic behaviour of consumers who exhibit no excess smoothness.

Near-rational behaviour

Drawing on the insight of Akerlof and Yellen (1985), Cochrane (1989) suggested that the excess sensitivity of consumption to expected income changes could be explained by "near-rational behaviour". The idea is that agents have bounded rationality and choose not to calculate the optimal consumption response to an income change when the change is small. This recalls the assumption of “rational inattention” because it implies some kind of “cost” faced by the consumer, even though related to reoptimization rather than to information processing, and a somehow “rational” choice of the consumer to bear that cost only when the associated welfare gain makes it worthwhile.

Cochrane (1989)’s calculations suggest that the utility loss from setting consumption equal to income instead of fully optimizing — i.e., the loss from rule-of-thumb behaviour — can be negligible. Accordingly, even though the observed behaviour deviates substantially from the predictions of an intertemporal utility maximization model, this actually implies small deviations in terms of welfare. The same insight underlies Souleles’s (1999), Hsieh’s (2003), and Scholnick’s (2013) contributions aimed at investigating whether the magnitude of income changes can explain the evidence of excess sensitivity. More recently, Kueng (2018) was able to use the properties of his data set and the features of the reported consumption behaviour in order to rule out most alternative explanations of excess sensitivity, including liquidity constraints, buffer-stock saving, imperfect information, and rational inattention. He calculated the utility loss encountered by households for deviating from utility-maximizing behaviour and found it to be extremely small (by appropriately smoothing the income change, households would gain less than a day of consumption per year), concluding that his findings are consistent with near-rational behaviour.

Temptation preferences

In two recent contributions, Attanasio, Kovacs, and Moran (2020 and 2021) build on the growing body of experimental evidence showing that households suffer from present bias that makes it difficult to accumulate wealth in liquid form (see, e.g., Thaler and Benartzi 2004).[26] Accordingly, they introduce the assumption of “temptation preferences” in a two-asset model of consumption behaviour in order to provide an explanation for the presence of wealthy hand-to-mouth consumers different from the one proposed by Kaplan and Violante (2014). The key feature of temptation preferences is that utility depends not only on actual consumption decisions, but also on the most tempting consumption alternative available in the choice set in each period. The idea is that households suffer from self-control problems and therefore find it difficult to save in liquid assets, due to the possibility of instantaneous gratification which is hard to resist. This creates a demand for commitment devices (and thus illiquidity) that allow households to mitigate self-control problems. Attanasio, Kovacs, and Moran (2020) show that temptation preferences can account for the large share of hand-to-mouth households reported in the data even when liquid assets are assumed to deliver higher returns than illiquid assets.

Mental accounting

One of the standard assumptions about consumers' preferences that is rejected by behavioural economics is that according to which resources are fungible. In fact, individuals tend to treat different sources of income and wealth in different ways, showing a greater willingness to spend certain sources than others. That allowing for such "mental accounting" could help explain the empirical evidence has been suggested from time to time in the Euler equation literature (by, e.g., Bernheim, Skinner, and Weinberg 2001 and Kueng 2018), but only a recent contribution by Baugh *et al.* (2021) provides an explicit test of the hypothesis. Baugh *et al.* (2021) find that the same households increase consumption when they receive expected tax refunds, but smooth consumption when making tax payments, i.e., households deviate from optimization when faced with anticipated income increases but behave as utility-maximizers when faced with anticipated income declines. This asymmetric consumption pattern appears inconsistent with liquidity constraints, buffer-stock saving, rational inattention, or hand-to-mouth behaviour and Baugh *et al.* (2021) explain it through a model of mental accounting (Shefrin and Thaler 1988). The idea is that households maintain three mental accounts: current income, current assets, and future income. Specifically, an anticipated tax refund is considered future income, which a rule of thumb prevents households from consuming before arrival; payments, on the other hand, are viewed as unrelated to the income-consumption process, so that they are funded from less liquid savings accounts and do not lead to consumption fluctuations.

SOME REMARKS ON THE CURRENT STATE OF CONSUMPTION RESEARCH

Our reconstruction shows the strong and indeed peculiar link between empirical research and theoretical contributions that characterizes the literature on consumption. As we have seen, the need to reconcile the analytical results with the data has constantly driven scholars to propose further extensions of the original model. For reasons of space, it has not been possible to discuss all the numerous versions of the model of intertemporal utility maximization available in the literature, but we hope to have nonetheless provided a sufficiently general overview of the developments that have taken place in consumption research over the last few decades.[27]

A preliminary remark we would like to point out stems from the acknowledgement that theoretical developments have essentially consisted in modifying some of the model's underlying assumptions with the aim of obtaining predictions that are compatible with observed consumption behaviour. Indeed, if we imagine to remove all the specific assumptions introduced in each extension, i.e., if we consider the model in its most general formulation in which both individual preferences and the budget constraint are not specified, then the theory of intertemporal choice boils down to the thesis according to which the consumer maximizes his lifetime utility on the basis of the (expected) resources available during the life cycle. Now, this only implies that (discounted) marginal utility should be constant over time. Since this in turn can involve quite different consumption profiles, the thesis on its own has no implications in terms of saving behaviour and cannot, therefore, be empirically verified or used to assess the impact of specific policies.[28] This, however, is not in itself necessarily a problem. Clearly, one can legitimately doubt whether the constant search for additional assumptions tailored on the evidence the model is supposed to replicate is the most rigorous way to endow a theory with microeconomic foundations. But even leaving aside this fundamental methodological problem, the fact is that each version of the model adopts a specific set of assumptions which, on the one hand, is crucial for explaining the empirical evidence taken into account and, on the other hand, leads to implications that are different from those obtainable with any other version of the model. But if this is the case, and if, at the same time, what is demanded from the theory of intertemporal choice is to describe the consumption behaviour actually observed and to provide policy prescriptions, then it becomes very much necessary to reach a general consensus on which particular version of the model is to be considered "the standard model". In other words, it is necessary to converge on the one hand on a specific definition of the stochastic environment in which the consumer operates, and therefore on a specific definition of the intertemporal budget constraint, and on the other hand on a precise specification of individual preferences, which includes both the choice of the utility function and the identification of the value to be assigned to preference parameters.[29] However, as witnessed by our reconstruction, more than four decades after the original formulation, no widely accepted version of the model seems to have emerged. On the contrary, the most recent years have seen the appearance of further new versions that aim to improve on the still not fully satisfactory empirical performance of previous models and that mark an increasingly radical departure from the original formulation proposed by Hall (1978).

In this respect, there is another issue we would like to address, which is related not so much to the *extent* as to the *nature* of such departure. To this end, it should first be recalled that as an early response to the problematic empirical results it was suggested that the model could incorporate a significant fraction of individuals who, instead of maximizing their utility, behave as rule-of-thumb consumers and simply set consumption equal to current income in each period. Clearly, following this route would have implied renouncing to the idea that the theory of intertemporal choice provides an acceptable representation of the observed consumption behaviour. And it is in fact for the purpose of escaping this outcome that scholars have proposed several extensions of the original model. They have in the first place modified the assumptions shaping the framework faced by the consumer and those regarding the value to be assigned to preference parameters. However, they have also increasingly relied on the introduction of elements that are somehow difficult to reconcile with the theoretical and methodological premises of the neoclassical approach. This is the case for the models in which individual utility depends on other consumers' behaviour, i.e., the models which allow for those social interdependencies in consumption that ever since the early stages in the development of neoclassical consumption theory were understood to have problematic analytical and methodological implications and were accordingly kept aside.[30] The same applies to those models that imply some form of boundary to the very concept of rationality underlying neoclassical theory, that is the models that regard the consumer as having a limited capacity of gathering or processing information or a limited capacity of identifying optimal behaviour. Only partially different is the question of the models that drop the hypothesis according to which preferences are additive over states of nature, since in this case it is not an apparently essential assumption of neoclassical theory that is violated, but still a cornerstone of the neoclassical approach to the study of choices under risk such as the expected-utility theory. However, what undoubtedly marks an authentic fracture is the recent incorporation of psychological elements borrowed from behavioural economics, which seems to contradict even the basic idea according to which individuals consistently maximize their utility. In this regard, it may be worthwhile to recall a passage of the review of the early Euler equation literature in which Browning and Crossley (2001) put forward a keen defence of the approach:

> In its most general formulation, the life-cycle framework simply asserts that agents make sequential decisions to achieve a coherent (and "stable") goal using currently available information as best they can. This catholic view does not rule out many models which would not be consistent with earlier restrictive models in the life-cycle tradition […]. What the life-cycle framework does rule out is "rule of thumb" behavior, in which households simply spend a fixed fraction of their income. It also rules out many psychological or behavioral explanations (pp. 3-4).

As a matter of fact, in the twenty years following this claim, the role of rule-of-thumb behaviour may have been mitigated but definitely not completely set aside, while those "behavioural explanations" have eventually been admitted into "the life-cycle framework". Overall, even leaving aside any more general criticism that could be addressed to the mainstream explanation of consumer behaviour, this process of progressive departure from the very premises of the approach suggests a substantial failure of the research project initiated with Hall's (1978) contribution and can arguably be seen as illustrative of that state of "fragmentation" (Roncaglia 2019) that characterizes contemporary economic theory.

Clearly, once the theory of intertemporal choice has been discarded, the question of the determinants of individuals' consumption behaviour remains open. In this regard, it is worth pointing out that the main results emerging from the empirical literature, relating to the role of current income as an important determinant of consumption ("excess sensitivity") and to the tendency of consumers not to change drastically their level of expenditure from one period to another ("excess smoothness"), are both accounted for when one adopts a conception of consumption as largely determined by social factors. The idea is that consuming certain goods allows individuals to be identified with specific social groups and that the concern for the display of their social status leads individuals to endeavour to maintain the level of consumption to which they are accustomed. As a consequence, when income decreases, consumers would rather let this affect their saving than compromise their acquired standard of living, which implies that consumption is to some extent "irreversible". By contrast, when income increases, individuals readily increase their consumption in order to flaunt a social status previously inaccessible to them. This could explain the asymmetrical responses of consumption to positive and negative changes in income that are observed in the data.[31]

It is important to emphasize that we do not mean to advocate the social role of consumption as the sole determinant of the observed behaviour, and in particular we would not deny the existence of a plan-making component in consumption choices that takes into account, among other things, individual expectations about future income. Rather, what the literature on consumption and saving behaviour seems to suggest is that the factors influencing consumption choices are too numerous and too complex to be adequately captured in a single model.

This leads us to our final remark. Because of the extensiveness of the literature addressed in the present paper, we have decided to devote a separate work to the policy implications of the various versions of the model of intertemporal choice proposed. However, one should never forget that in the background of the competition between the "Keynesian" consumption function and the alternatives proposed during the debate of the 1940s and 1950s there was the question of the effectiveness of fiscal policy. In this regard, it should be noted that the empirical results of the Euler equation literature do not seem to support Friedman's thesis based on the irrelevance of current income as a determinant of current consumption — and even less that "Ricardian" equivalence which takes such thesis to its extreme consequences. As a matter of fact, within the mainstream literature itself, it has become common practice to introduce into dynamic general equilibrium models the hypothesis of a fraction of consumers who behave as "non-Ricardians" (either because they are liquidity constrained or simply because they are rule-of-thumb consumers) as a channel through which fiscal policy can have "Keynesian" effects.

Acknowledgment

While working on this paper, I benefited from the help of many people. I would like to thank in particular Antonella Palumbo for the valuable suggestions she gave me at differents times and Attilio Trezzini for introducing me to the economics of consumption and for his constant encouragement and support. I am also grateful to two anonymous referees of this journal for helpful comments.

Notes

1. The distinction between transitory and permanent income is able to explain the difference between the long-run and the short-run marginal propensity to consume. According to Friedman, the level of income observed in cross-section data is likely to be affected by transitory components, which leads to a lower estimate of the marginal propensity to consume; as the length of the

period of observation increases, the influence of transitory components gets smaller and "measured" income becomes very close to permanent income, so that, in time-series data, consumption is found to be proportional to income. There were other well-known "stylized facts" that were calling for an explanation. For example, the evidence suggested that black families consume a larger proportion of their income than white families, but, at the same level of income, black families consume less than white families. According to the LCPIH, this happens because blacks tend to have lower permanent incomes than whites and, therefore, at any given level of *current* income, they tend to have larger *transitory* income and hence save more.

2. Another influential theory proposed as an explanation for the puzzling empirical evidence at the centre of the debate on the consumption function is the relative income hypothesis developed by Duesenberry (1949). In his contribution, Friedman (1957) recognizes that in terms of empirical performance Duesenberry's theory is just as successful as the permanent income hypothesis, but advocates his own approach as being more internally consistent. For a survey of the debate on the consumption function which includes a comparison between Duesenberry's and Friedman's theories, see Trezzini (2012).
3. Following the literature, we will generally identify the consumer with a generic "individual", but the decision-making unit is usually assumed to be the household.
4. Actually, Hall (1978) admitted that "the pure life cycle-permanent income hypothesis […] is rejected by the data" and suggested a modification of the hypothesis according to which "some part of consumption takes time to adjust to a change in permanent income". He claimed the data to be "entirely compatible" with this modification of the hypothesis, which, according to him, "leaves its central content unchanged" (p. 985). However, this aspect of Hall's contribution has remained virtually unnoticed.
5. To the best of my knowledge, there is no single work providing an exhaustive and accurate discussion of the assumptions underlying the Euler equation approach to consumption behaviour. There are, however, contributions that partially serve this purpose, such as Muellbauer (1983), Browning and Lusardi (1996), and Attanasio (1999).
6. Notice that quadratic utility actually violates the monotonicity assumption, as it implies a bliss level for consumption. However, the assumption can be weakened to state merely that no consumption level attainable with the individual's resources is preferred to all others.
7. Campbell (1987) derives an expression for consumption as a function of the expected value of lifetime resources under the assumption of quadratic preferences. Caballero (1990) addresses the case of exponential utility. Both preference specifications, however, are not very appealing (see, e.g., Jappelli and Pistaferri 2017).

8. The model is often referred to as the “permanent-income hypothesis” or the “life-cycle model”, emphasizing the continuity with the theories of Modigliani and Friedman. However, the formal treatment of uncertainty introduced with the assumption of rational expectations seems to have brought about a marked change with respect to the analyses of the 1950s. The expression “standard model” is also very popular, but for reasons that will become clear later, we find it inappropriate as well. We shall therefore speak of the “Euler equation approach” or the “intertemporal (utility) maximization model”.
9. At the empirical level, excess sensitivity arises when the data show a positive correlation between the change in consumption and the lagged change in income.
10. Technically, excess smoothness emerges because the empirical variance of consumption innovations is smaller than the variance of innovations in permanent income.
11. Jappelli and Pistaferri (2017) describe the strategies used to identify anticipated and unanticipated income changes and the empirical techniques adopted in excess sensitivity and excess smoothness tests, providing also a survey of the main findings. Fuchs-Schündeln and Hassan’s (2016) review of the literature on “natural experiments” also focuses on both the methodological side and the empirical results. Finally, it is worth mentioning two recent meta-analyses concerning the excess smoothness (Havránek, Rusnák, and Sokolova 2017) and the excess sensitivity (Havránek and Sokolova 2020) puzzles.
12. Browning and Crossley (2001) provide a discussion of the evidence of consumption smoothing at different frequencies: high (within the year), medium (year to year or across the business cycle), low (across the working life), and very low (across stages of the life cycle).
13. More recently, similar remarks can be found in Jappelli and Pistaferri (2017) (see, e.g., p. 28). Cohen *et al.* (2020) review the research that measures time preferences. For an overview of the evidence on the magnitude of risk aversion, see Ludvigson (2013) and Outreville (2014). Browning and Lusardi (1996) provide a detailed survey of the early contributions on precautionary saving, and Lugilde, Bande, and Riveiro (2019) review the recent evidence. As for the elasticity of intertemporal substitution, detailed accounts are provided in the meta-analysis by Havránek *et al.* (2015) and the survey by Thimme (2017).
14. For a more extensive discussion of the analytical definition of preference parameters and of the controversial results of the literature devoted to their estimation, see Pignalosa (2019).
15. Both theoretical (Arrow 1965) and empirical arguments suggest that the value of the coefficient of relative risk aversion should be equal to one or at least not much higher. However, an example may help appreciate the implausibility of the values reported by Mehra and Prescott (1985). Let us then consider a lottery which yields, with equal probability, 50 or 100 units of consumption. An individual who is neither risk averse nor risk loving, in order to be willing to give up playing the lottery, should receive a sure amount of (at least) 75

units of consumption (the expected value of the lottery). Let us now consider a risk averse individual with standard isoelastic preferences. To make him willing to forego the lottery, we should offer him 70.7 units of consumption if his coefficient of relative risk aversion is equal to 1, 54.0 units if the coefficient is 10, and only 51.9 units if the coefficient is 20.

16. Strictly speaking, consumers are “rule-of-thumb” if, rather than behaving as utility-maximizers, they follow simple or even relatively sophisticated heuristics. However, in the mainstream literature, the term “rule-of-thumb” is usually used in a narrower sense, as denoting those individuals who follow the particular rule which prescribes to consume all of their income in every period.
17. As a matter of fact, all consumers who simply consume their current income, regardless of the reason, are “hand-to-mouth”. Their behaviour could either be the result of a rule of thumb or, for example, “the reflection of an inability to trade in asset markets due to infinite transactions costs” (Weil 1992: p. 575). However, the term is frequently used to refer specifically to non-optimizing behaviour.
18. They are referred to as “non-Ricardian” consumers because the so-called “Ricardian equivalence” (Barro 1974) does not hold for households consuming their current income period by period.
19. In particular, the theory suggests that a consistently aggregated Euler equation would be based on *the sum of the logarithms* of individual consumption, but what can be calculated from national accounts is *the logarithm of the sum* of individual consumption. Attanasio and Weber (1993) show the relevance of such aggregation bias and suggest that consistently aggregated micro data partially solve rejections of the model that instead appear by making use of macroeconomic data.
20. A discussion of these more technical aspects is beyond our scope; the interested reader is referred to Jappelli and Pistaferri (2017).
21. Consider the quadratic utility function adopted by Hall (1978). Under this preference specification, the marginal utility of current consumption is $\bar{C} - C_t$ and, because of the linearity of the expectation operator, the expected value of the marginal utility of next-period consumption is $\bar{C} - \mathbb{E}_t[C_{t+1}]$, which is the same as the marginal utility of the expected value of next-period consumption: $\mathbb{E}_t[u'(C_{t+1})] = u'(\mathbb{E}_t[C_{t+1}])$. Since it depends only on expected consumption, expected marginal utility is not affected by a mean-preserving increase in uncertainty. When preferences allow for precautionary saving, instead, the Euler equation also includes the conditional variance of consumption growth.
22. If marginal utility is convex, then, by definition, $\mathbb{E}_t[u'(C_{t+1})] > u'(\mathbb{E}_t[C_{t+1}])$, i.e., the expected marginal utility of consumption is larger than the marginal utility of expected consumption. This implies that a mean-preserving increase in uncertainty raises expected marginal utility. In order to bring expected marginal utility back into equality with current marginal utility, current consumption must be reduced.

23. For an analysis of the interplay of liquidity constraints and prudence see Carroll and Kimball (2001) and Carroll, Holm, and Kimball (2021).
24. In the intertemporal maximization model, lifetime utility is the sum of utilities from consumption in each period, where utility of future consumption is expected utility, i.e., the sum of utilities in each state of nature weighted with the probabilities of each state. Von Neumann and Morgenstern's (1947) "expected utility theorem" shows that if preferences satisfy the completeness, transitivity, continuity, and independence axioms, then they admit a utility representation that is linear in probabilities, which implies that the utility of an aleatory amount can be thought of as the expected value of the utilities of the different outcomes. As is well known, the axioms underlying the expected-utility representation of preferences are quite controversial. For a reconstruction of the process that led to their acceptance and the establishment of expected-utility theory as the mainstream approach to the analysis of risky choices, see Moscati (2016).
25. For a thorough discussion of EZW preferences see Kimball and Weil (2009).
26. The idea is that in discounting future utility with respect to present utility, the consumer is affected by a bias that leads him to significantly reduce the weight attached to future utility.
27. Actually, we have not discussed all the empirical puzzles either. For example, the model of intertemporal utility maximization struggles to explain differences in saving rates across countries. As for the extensions of the model, we have left out, for example, the formulations in which consumers have bequest motives or face uncertainty about the length of life; we have also neglected the literature analysing hyperbolic discounting and that focusing on the relevance of financial sophistication. For an overview of all these strands of research, see Jappelli and Pistaferri (2017).
28. As suggested by Browning and Lusardi (1996: p. 1800), "the most general model […] does not seem to impose any restrictions on the time path of consumption and asset prices. It is only when we impose restrictions on preferences and budgets that we can derive testable implications. Thus the standard model in its most general form is better thought of as a framework than as a direct source of testable propositions". Similar statements are to be found in Browning and Crossley (2001: p. 3) and in Attanasio and Weber (2010: p. 695).
29. Clearly, in order to deliver precise *quantitative* predictions, any theory requires some parameters to be given specific values. However, the model of intertemporal utility maximization is unable to provide even *qualitative* indications about consumption behaviour unless all preference parameters are quite precisely set. The most patent example of this comes from the comparison between the model with quadratic preferences and the model with precautionary saving. In the first one the consumer is supposed to plan for a constant consumption level throughout the life cycle (and consumption is supposed to change only if unexpected income changes occur); by contrast,

a precautionary motive for saving implies a growing trend of consumption. The difference is entirely due to a single parameter — prudence — being either null or positive. And the theory of intertemporal choice provides no a priori reason for expecting prudence to behave one way or the other, so that the relevance of the precautionary motive for saving is a purely empirical matter.

30. The main problems arise in connection to the construction of a market demand function as the summation of individual demands and to the measurement of consumer surplus. On this point, see Drakopulos (2012) and Bianchi and Sanfilippo (2015).

31. The emergence of an approach which regards consumption as a primarily social phenomenon can be traced to the work of Veblen (1899). Veblen's insights were extensively used within the debate on the consumption function of the 1940s and 1950s. Besides Duesenberry's (1949) contribution, it is worth mentioning the work of Samuelson (1943) and a contribution by Modigliani (1949) that preceded his much more successful life-cycle model (for a reconstruction, see Trezzini 2005). They have also been introduced in the Euler equation literature in the form of habit-forming preferences and in the "keeping/catching up with the Joneses" models, but the social significance of consumption does not seem easily reconcilable with the neoclassical approach.

References

Abel, A.B. (1990), "Asset prices under habit formation and catching up with the Joneses", *American Economic Review*, 80(2), pp. 38-42.

Aguiar, M. and Hurst, E. (2005), "Consumption versus expenditure," *Journal of Political Economy*, 113(5), pp. 919-948.

Akerlof, G.A. and Yellen, J.L. (1985), "A near-rational model of the business cycle, with wage and price inertia", *Quarterly Journal of Economics*, 100(Supplement), pp. 823-838.

Arrow, K.J. (1965), *Aspects of the Theory of Risk Bearing*, Helsinki, Finland, Yrjo Jahnsson Lectures.

Attanasio, O.P. (1995), "The intertemporal allocation of consumption: theory and evidence", *Carnegie-Rochester Conference Series on Public Policy*, 42, pp. 39-89.

Attanasio, O.P. (1999), "Consumption", in Taylor, J.B. and Woodford, M. (eds), *Handbook of Macroeconomics*, volume 1, Amsterdam, Elsevier, pp. 741-812.

Attanasio, O.P. and Browning, M. (1995), "Consumption over the life cycle and over the business cycle", *American Economic Review*, 85(5), pp. 1118-1137.

Attanasio, O.P. and Weber, G. (1989), "Intertemporal substitution, risk aversion and the Euler equation for consumption", *Economic Journal*, 99(395), pp. 59-73.

Attanasio, O.P. and Weber, G. (1993), "Consumption growth, the interest rate and aggregation", *Review of Economic Studies*, 60(3), pp. 631-649.

Attanasio, O.P. and Weber, G. (1995), "Is consumption growth consistent with intertemporal optimization? Evidence from the consumer expenditure survey", *Journal of Political Economy*, 103(6), pp. 1121-1157.

Attanasio, O.P. and Weber, G. (2010), "Consumption and saving: models of intertemporal allocation and their implications for public policy", *Journal of Economic Literature*, 48(3), pp. 693-751.

Attanasio, O.P., Banks, J., Meghir, C., and Weber, G. (1999), "Humps and bumps in lifetime consumption", *Journal of Business and Economic Statistics*, 17(1), pp. 22-35.

Attanasio, O.P., Kovacs, A., and Moran, P. (2020), "Temptation and commitment: understanding hand-to-mouth behavior", *NBER Working Papers*, 27944.

Attanasio, O.P., Kovacs, A., and Moran, P. (2021), "Temptation and incentives to wealth accumulation", *NBER Working Papers*, 28938.

Bansal R. and Yaron, A. (2004), "Risks for the long run: a potential resolution of asset pricing puzzles", *Journal of Finance*, 59(4), pp. 1481-1509.

Barro, R.J. (1974), "Are Government bonds net wealth?", *Journal of Political Economy*, 82(6), pp. 1095-1117.

Baugh, B., Ben-David, I., Park, H., and Parker, A.J. (2021), "Asymmetric consumption smoothing", *American Economic Review*, 111(1), pp. 192-230.

Bernheim, B.D. (1987), "Dissaving after retirement: testing the pure life cycle hypothesis", in Bodie, Z., Shoven, J.B., and Wise, D.A. (eds), *Issues in Pension Economics*, University of Chicago Press, pp. 237-280.

Bernheim, B.D., Skinner, J., and Weinberg, S. (2001), "What accounts for the variation in retirement wealth among U.S. households?", *American Economic Review*, 91(4), pp. 832-857.

Bianchi, M. and Sanfilippo, E. (2015), "Social interdependencies in consumption: an early economic debate on social distinction, emulation and fashion", *History of Political Economy*, 47(4), pp. 547-575.

Blanchard, O.J. and Mankiw, N.G. (1988), "Consumption: beyond certainty equivalence", *American Economic Review*, 78(2), pp. 173-177.

Blundell, R., Browning, M., and Meghir, C. (1994), "Consumer demand and the life-cycle allocation of household expenditure", *Review of Economic Studies*, 61(1), pp. 57-80.

Breeden, D.T. (1979), "An intertemporal capital asset pricing model with stochastic consumption and investment opportunities", *Journal of Financial Economics*, 7(3), pp. 265-296.

Breeden, D.T., Gibbons, X., and Litzenberger, R. (1989), "Empirical tests of the consumption-oriented CAPM", *Journal of Finance*, 44(2), pp. 231-262.

Browning, M. and Crossley, A. (2001), "The life-cycle model of consumption and saving", *The Journal of Economic Perspectives*, 15(3), pp. 3-22.

Browning, M. and Lusardi, A. (1996), "Household saving: micro theories and macro facts", *Journal of Economic Literature*, 34(4), pp. 1797-1855.

Browning, M., Deaton, A., and Irish, M. (1985), "A profitable approach to labor supply and commodity demands over the life-cycle", *Econometrica*, 53(3), pp. 503-543.

Caballero, R.J. (1990), "Consumption puzzles and precautionary savings", *Journal of Monetary Economics*, 25(1), pp. 113-136.

Campbell, J.Y. (1987), "Does saving anticipate declining labor income? An alternative test of the permanent income hypothesis", *Econometrica*, 55(6), pp. 1249-1273.

Campbell, J.Y. and Cochrane, J.H. (1999), "By force of habit: a consumption-based explanation of aggregate stock market behavior", *Journal of Political Economy*, 107(2), pp. 205-251.

Campbell, J.Y. and Deaton, A.S. (1989), "Why is consumption so smooth?", *Review of Economic Studies*, 56(3), pp. 357-373.

Campbell, J.Y. and Mankiw, N.G. (1989), "Consumption, income, and interest rates: reinterpreting the time series evidence, *NBER Macroeconomics Annual*, 4, pp. 185-216.

Campbell, J.Y. and Mankiw, N.G. (1991), "The response of consumption to income: a cross-country investigation", *European Economic Review*, 35(4), pp. 723-767.

Carroll, C.D. (1992), "The buffer-stock theory of saving: some macroeconomic evidence", *Brookings Papers on Economic Activity*, 23(2), pp. 61-156.

Carroll, C.D. (1997), "Buffer-stock saving and the life cycle/permanent income hypothesis", *Quarterly Journal of Economics*, 112(1), pp. 1-55.

Carroll, C.D. (2001), "Death to the log-linearized consumption Euler equation! (And very poor health to the second-order approximation)", *The B.E. Journal of Macroeconomics*, 1, pp. 1-18.

Carroll, C.D. (2011), "Theoretical foundations of buffer stock saving", *CFS Working Papers*, 2011/15.

Carroll, C.D. and Kimball, M.S. (2001), "Liquidity constraints and precautionary saving", *NBER Working Papers*, 8496.

Carroll, C.D. and Summers, L.H. (1991), "Consumption growth parallels income growth: some new evidence", in Bernheim, B.D. and Shoven, J.B. (eds), *National Saving and Economic Performance*, Chicago: University of Chicago Press, pp. 305-348.

Carroll, C.D., Crawley, E., Slacalek, J., Tokuoka, K., and White, M.N. (2020), "Sticky expectations and consumption dynamics", *American Economic Journal: Macroeconomics*, 12(3), pp. 40-76.

Carroll, C.D., Holm, M.B., and Kimball, M.S. (2021), "Liquidity constraints and precautionary saving", *Journal of Economic Theory*,195(C), pp. 1-21.

Cochrane, J.H. (1989), "The sensitivity of tests of the intertemporal allocation of consumption to near-rational alternatives", *American Economic Review*, 79(3), pp. 319-337.

Cohen, J.D., Ericson, K.M., Laibson, D., and White, J.M. (2020), "Measuring time preferences", *Journal of Economic Literature*, 58(2), pp. 299-347.

Constantinides, G.M. (1990), "Habit formation: a resolution of the equity premium puzzle", *Journal of Political Economy*, 98(3), pp. 519-543.

De Giorgi, G., Frederiksen, A., and Pistaferri, L. (2020), "Consumption network effects", *Review of Economic Studies*, 87(1), pp. 130-163.

Deaton, A.S. (1986), "Life-cycle models of consumption: is the evidence consistent with the theory?", *NBER Working Papers*, 1910.

Deaton, A.S. (1991), "Saving and liquidity constraints", *Econometrica*, 59(5), pp. 1221-1248.

Deaton, A.S. (1992): *Understanding Consumption*, Oxford: Oxford University Press.

Diamond, P.A. and Hausman, J.A. (1984), "Individual retirement and savings behavior", *Journal of Public Economics*, 23(1-2), pp. 81-114.

Drakopulos, S.A. (2012), "The history of attitudes towards interdependent preferences", *Journal of the History of Economic Thought*, 34(4), pp. 541-557.

Duesenberry, J.S. (1949), *Income, Saving and the Theory of Consumption Behavior*, Cambridge, Massachusetts: Harvard University Press.

Dynan, K.E. (1993), "How prudent are consumers?", *Journal of Political Economy*, 101(6), pp. 1104-1113.

Dynan, K.E. (2000), "Habit formation in consumer preferences: evidence from panel data", *American Economic Review*, 90(3), pp. 391-406.

Eisenhauer, J.G. (2000), "Estimating prudence", *Eastern Economic Journal*, 26(4), pp. 379-392.

Epstein, L.G. and Zin, S.E. (1989), "Substitution, risk aversion, and the temporal behavior of consumption and asset returns: a theoretical framework", *Econometrica*, 57(4), pp. 937-969.

Epstein, L.G. and Zin, S.E. (1991), "Substitution, risk aversion, and the temporal behavior of consumption and asset returns: an empirical analysis", *Journal of Political Economy*, 99(2), pp. 263-286.

Flavin, M. (1981), "The adjustment of consumption to changing expectations about future income", *Journal of Political Economy*, 89(5), pp. 974-1009.

Friedman, M. (1957), *A Theory of the Consumption Function*, Princeton: Princeton University Press.

Fuchs-Schündeln, N. and Hassan, T.A. (2016), "Natural experiments in macroeconomics", in Taylor, J.B. and Uhlig, H. (eds), *Handbook of Macroeconomics*, volume 2, Amsterdam, Elsevier, pp. 923-1012.

Fuhrer, J.C. (2000), "Habit formation in consumption and its implications for monetary-policy models", *American Economic Review*, 90(3), pp. 367-390.

Galí, J., López-Salido, J.D., and Vallés, J. (2004), "Rule-of- thumb consumers and the design of interest rate rules", *Journal of Money, Credit and Banking*, 36(4), pp. 739-764.

Grossman, S., Melino, A., and Shiller, R. (1987), "Estimating the continuous-time consumption-based asset-pricing model", *Journal of Business and Economic Statistics*, 5(3), pp. 315-327.

Guiso, L., Jappelli, T., and Terlizzese, D. (1992), "Earnings uncertainty and precautionary saving", *Journal of Monetary Economics*, 30(2), pp. 307-337.

Hall, R.E. (1978), "Stochastic implications of the life cycle-permanent income hypothesis: theory and evidence", *Journal of Political Economy*, 86(6), pp. 971-987.

Hall, R.E. (1988), "Intertemporal substitution in consumption", *Journal of Political Economy*, 96(2), pp. 339-357.

Hamermesh, D.S. (1984), "Consumption during retirement: the missing link in the life cycle", *Review of Economics and Statistics*, 66(1), pp. 1-7.

Hansen, L.P. and Singleton, K.J. (1983), "Stochastic consumption, risk aversion, and the temporal behavior of asset returns", *Journal of Political Economy*, 91(2), pp. 249-268.

Havránek, T. and Sokolova, A. (2020), "Do consumers really follow a rule of thumb? Three thousand estimates from 144 studies say "probably not"", *Review of Economic Dynamics*, 35, pp. 97-122.

Havránek, T., Horvath, R., Irsova, Z., and Rusnák, M. (2015), "Cross-country heterogeneity in intertemporal substitution", *Journal of International Economics*, 96(1), pp. 100-118.

Havránek, T., Rusnák, M., and Sokolova, A. (2017), "Habit formation in consumption: a meta-analysis", *European Economic Review*, 95(C), pp. 142-167.

Hayashi, F. (1985), "The effect of liquidity constraints on consumption: a cross-sectional analysis, *Quarterly Journal of Economics*, 100(1), pp. 183-206.

Hsieh, C.-T. (2003), "Do consumers react to anticipated income changes? Evidence from the Alaska Permanent Fund", *American Economic Review*, 93(1), pp. 397-405.

Jappelli, T. and Pistaferri, L. (2010), "The consumption response to income changes", *Annual Review of Economics*, 2(1), pp. 479-506.

Jappelli, T. and Pistaferri, L. (2017), *The Economics of Consumption. Theory and Evidence*, Oxford: Oxford University Press.

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

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

Kaplan, G.W., Violante, G.L., and Weidner, J. (2014), "The wealthy hand-to-mouth", *Brookings Papers on Economic Activity*, 45(1), pp. 77-138.

Keynes, J.M. (1936), *The General Theory of Employment, Interest and Money*, Cambridge: Macmillan.

Kiley, M.T. (2010), "Habit persistence, nonseparability between consumption and leisure, or rule-of-thumb consumers: which accounts for the predictability of consumption growth?", *Review of Economics and Statistics*, 92(3), pp. 679-683.

Kimball, M.S. and Weil, P. (2009), "Precautionary saving and consumption smoothing across time and possibilities", *Journal of Money, Credit and Banking*, 41(2), pp. 245-284.

Kreps, D.M. and Porteus, E.L. (1978), "Temporal resolution of uncertainty and dynamic choice theory", *Econometrica*, 46(1), pp. 185-200.

Kueng, L. (2018), "Excess sensitivity of high-income consumers", *Quarterly Journal of Economics*, 133(4), pp. 1693-1751.

Lettau, M. and Uhlig, H. (1999), "Rules of thumb versus dynamic programming", *American Economic Review*, 89(1), pp. 148-174.

Ljungqvist, L. and Uhlig, H. (2000), "Tax policy and aggregate demand management under catching up with the Joneses", *American Economic Review*, 90(3), pp. 356-366.

Lucas, R.E. (1973)," Some international evidence on output-inflation trade-offs", *American Economic Review*, 63(3), pp. 326-334.

Lucas, R.E. (1976), "Econometric policy evaluation: a critique", *Carnegie-Rochester Conference Series on Public Policy*, 1, pp. 19-46.

Lucas, R.E. (1978), "Asset prices in an exchange economy", *Econometrica*, 46(6), pp. 1429-1445.

Ludvigson, S.C. (2013), "Advances in consumption-based asset pricing: empirical tests", in Constantinides, G.M., Harris, M., and Stulz, R. (eds), *Handbook of the Economics of Finance*, volume 2, Amsterdam, Elsevier, pp. 799-906.

Ludvigson, S.C. and Michaelides, A. (2001), "Does buffer-stock saving explain the smoothness and excess sensitivity of consumption?", *American Economic Review*, 91(3), pp. 631-647.

Lugilde, A., Bande, R., and Riveiro, D. (2019), "Precautionary saving: a review of the empirical literature", *Journal of Economic Surveys*, 33(2), pp. 481-515.

Mankiw, N.G. (2000), "The savers–spenders theory of fiscal policy", *American Economic Review*, 90(2), pp. 120-125.

Meghir, C. and Weber, G. (1996), "Intertemporal non-separability or borrowing restrictions? A disaggregate analysis using a U.S. consumption panel", *Econometrica*, 64(5), pp. 1151-1181.

Mehra, R. and Prescott, E.C. (1985), "The equity premium: a puzzle", *Journal of Monetary Economics*, 15(2), pp. 145-161.

Mirer, T.W. (1979), "The wealth-age relation among the aged", *American Economic Review*, 69(3), pp. 435-443.

Modigliani, F. (1949), "Fluctuations in the saving-income ratio: a problem in economic forecasting", *Studies in Income and Wealth*, volume 11, NBER, pp. 371-441.

Modigliani, F. and Brumberg, R. (1954), "Utility analysis and the consumption function: an interpretation of cross-section data", in Kurihara, K. (ed.), *Post Keynesian Economics*, New Brunswick: Rutgers University Press.

Moscati, I. (2016), "Retrospectives: how economists came to accept expected utility theory: the case of Samuelson and Savage", *Journal of Economic Perspectives*, 30(2), pp. 219-236.

Muellbauer, J. (1983), "Surprises in the consumption function", *Economic Journal*, 93(Supplement), pp. 34-50.

Neumann, J. von and Morgenstern, O. (1947), *Theory of Games and Economic Behavior*, Princeton: Princeton University Press.

Outreville, J.F. (2014), "Risk aversion, risk behavior, and demand for insurance: a survey", *Journal of Insurance Issues*, 37(2), pp. 158-186.

Pignalosa, D. (2019), "On the role of the utility function in the estimation of preference parameters", *Metroeconomica*, 70(4), pp. 793-820.

Pischke, J.-S. (1995), "Individual income, incomplete information, and aggregate consumption", *Econometrica*, 63(4), pp. 805-840.

Roncaglia, A. (2019), *The Age of Fragmentation: A History of Contemporary Economic Thought*, Cambridge: Cambridge University Press.

Rubinstein, M. (1976), "The valuation of uncertain income streams and the pricing of options", *Bell Journal of Economics*, 7(2), pp. 407-425.

Samuelson, P.A. (1943), "Full employment after the war", in Harris, S.E. (ed.), *Postwar Economic Problems*, New York, McGraw Hill, pp. 27-53.

Scholnick, B. (2013), "Consumption smoothing after the final mortgage payment: testing the magnitude hypothesis", *Review of Economics and Statistics*, 95(4), pp. 1444-1449.

Shefrin, H.M. and Thaler, R.H. (1988), "The behavioral life-cycle hypothesis", *Economic Inquiry*, 26(4), pp. 609-643.

Sims, C.A. (2003), "Implications of rational inattention", *Journal of Monetary Economics*, 50(3), pp. 665-690.

Souleles, N.S. (1999), "The response of household consumption to income tax refunds", *American Economic Review*, 89(4), pp. 947-958.

Thaler, R.H. and Benartzi, S. (2004), "Save more tomorrow: using behavioral economics to increase employee saving", *Journal of Political Economy*, 102(S1), pp. 164-187.

Thimme, J. (2017), "Intertemporal substitution in consumption: a literature review", *Journal of Economic Surveys*, 31(1), pp. 226-257.

Thurow, L. (1969), "The optimum lifetime distribution of consumption expenditures", *American Economic Review*, 59(3), pp. 324-330.

Trezzini, A. (2005), "The economics of consumption as a social phenomenon: a neglected approach to the analysis of consumption", *Quaderni di Ricerca* *del Dipartimento Innovazione e Società*, Università degli Studi di Roma "La Sapienza", n. 2.

Trezzini, A. (2012), "Relative income vs. permanent income: the crisis of the theory of the social significance of consumption", *Journal of the History of Economic Thought*, 34(3), pp. 355-377.

Veblen, T. (1899), *The Theory of the Leisure Class: An Economic Study of Institutions*, New York, Macmillan.

Weil, P. (1989), "The equity premium puzzle and the risk-free rate puzzle", *Journal of Monetary Economics,* 24(3), pp. 401-421.

Weil, P. (1992), "Hand-to-mouth consumers and asset prices", *European Economic Review*, 36(2-3), pp. 575-583.

West, K.D. (1988), "The insensitivity of consumption to news about income", *Journal of Monetary Economics*, 21(1), pp. 17-33.

Winter, J.K., Schlafmann, K., and Rodepeter, R. (2012), "Rules of thumb in life-cycle saving decisions", *Economic Journal*, 122(560), pp. 479-501.

Zeldes, S.P. (1989a), "Consumption and liquidity constraints: an empirical investigation", *Journal of Political Economy*, 97(2), pp. 305-346.

Zeldes, S.P. (1989b), "Optimal consumption with stochastic income: deviations from certainty equivalence", *Quarterly Journal of Economics*, 104(2), pp. 275-298.