

Maurer, Jürgen; Meier, André

Working Paper

Do the Joneses really matter? Peer-group versus correlated effects in intertemporal consumption choice

IFS Working Papers, No. 05/15

Provided in Cooperation with:

The Institute for Fiscal Studies (IFS), London

Suggested Citation: Maurer, Jürgen; Meier, André (2005) : Do the Joneses really matter? Peer-group versus correlated effects in intertemporal consumption choice, IFS Working Papers, No. 05/15, Institute for Fiscal Studies (IFS), London, <https://doi.org/10.1920/wp.ifs.2005.0515>

This Version is available at:

<https://hdl.handle.net/10419/71506>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

DO THE “JONESES” REALLY MATTER? PEER-
GROUP VERSUS CORRELATED EFFECTS IN
INTERTEMPORAL CONSUMPTION CHOICE

Jürgen Maurer
André Meier

Do the "Joneses" Really Matter? Peer-group vs. Correlated Effects in Intertemporal Consumption Choice

Jürgen Maurer* and André Meier†

First version: June 2003

This version: July 2005

Abstract

Recent theoretical contributions have suggested consumption externalities, or peer-group effects, as a potential explanation for some of the puzzles in macroeconomics and finance. However, the empirical relevance of peer effects for intertemporal consumption choice is a completely open question. To shed some light on the issue, we derive an extension of the standard life-cycle model that allows for consumption externalities. The analysis is complicated by the challenge of disentangling actual peer effects from merely correlated effects operating through common features or shocks within peer groups. We show how to conduct reliable inference under these circumstances based on within-group equilibrium conditions that give rise to a social multiplier. This approach can be understood as an adaptation of Manski's "reflection problem framework" to the case of dynamic models with endogenous regressors. We estimate our model using US panel data from the PSID. While there is strong predictable consumption co-movement within peer groups, the evidence for true consumption externalities vanishes once correlated effects are adequately accounted for.

JEL classification: C23, D12, D91, Z13

Keywords: Consumption, Life-Cycle Model, Peer Effects, Reflection Problem

*European University Institute, Florence, Italy and Institute for Fiscal Studies, London, United Kingdom. Maurer gratefully acknowledges financial support provided by the European Community's Human Potential Programme through contract HPRN-CT-2002-00235, [AGE]. Address for correspondence: Jürgen Maurer, Institute for Fiscal Studies, 7 Ridgmount Street, London WC1E 7AE, United Kingdom. Email: juergen.maurer@iue.it.

†European University Institute, Florence, Italy. Email: andre.meier@iue.it.

We would like to thank Giuseppe Bertola, Charles Grant, Roger Klein and Frank Vella for very helpful discussion. Moreover, we have received valuable comments from Orazio Attanasio, Richard Blundell, Steven Bond, Martin Browning, Tullio Jappelli, Søren Leth-Petersen, Charles Manski, Robert Waldman, Guglielmo Weber, Frank Windmeijer and seminar participants at the 3rd Villa Mondragone Workshop in Rome, the EEA Summer School at the Cemmap in London, the Zeuthen Workshop at the CAM Copenhagen, the Institute for Fiscal Studies in London, the Central European University in Budapest and the University of Bern. Needless to say, all remaining errors are our own responsibility.

Executive Summary

Recent theoretical contributions have suggested consumption externalities, or peer-group effects, as a potential explanation for some of the puzzles in macroeconomics and finance. Despite its intuitive appeal, however, the idea of peer effects in intertemporal consumption choice has not been put to the empirical test.

To shed some light on the issue, we derive an extension of the standard life-cycle model that allows for consumption externalities. Specifically, we focus on a "keeping up with the Joneses" specification of individual utility. Our question is whether the choice of optimal consumption profiles is affected by the simultaneous consumption decisions of households with similar characteristics. The analysis is complicated by what is known as the "reflection problem", i.e. the challenge of disentangling actual peer effects from merely correlated effects operating through common features or shocks within peer groups. We show how to conduct reliable inference under these circumstances based on within-group equilibrium conditions that give rise to a social multiplier. This approach can be understood as an adaptation of Manski's reflection problem framework to the case of dynamic models with endogenous regressors.

We estimate our model using US panel data from the PSID. Reference groups are constructed on the basis of age, education, gender, race and urbanity. Although our results show strong predictable consumption co-movement within such reference groups, the evidence for true consumption externalities vanishes once correlated effects are adequately accounted for. Thus our study does not lend support to the widespread use of theoretical specifications assuming pronounced peer effects in intertemporal consumption choice.

1 Introduction

Consumption is arguably a social experience, and the position of other people with respect to our own consumption may often matter to us. This is reflected, for example, in notions like "conspicuous consumption", "peer-group effects" or "keeping up with the Joneses," which are commonplace in casual discussions about the determinants of particular consumption patterns. In line with this, psychologists, sociologists, and economists have collected evidence showing that consumers' well-being is affected by their *relative* economic standing rather than their absolute resources alone.¹ Early discussions of consumption externalities in economics date back at least to the seminal contributions of Duesenberry (1949) and Leibenstein (1950). It is thus very surprising that economists have largely ignored the issue when modelling intertemporal consumption choice.

Indeed, some "stylized facts" about life-cycle consumption seem suggestive of or at least compatible with peer effects. Rather than being smooth, life-cycle profiles of household consumption feature humps and bumps whose exact shapes appear to depend on characteristics of the respective households. Clearly this could be a consequence of changing demographic situations that have a direct effect on the utility derived from a given level of consumption.² However, we would expect to observe similarly synchronized consumption patterns within relevant peer groups if consumption externalities are important. Given the empirical support for peer-group effects in other fields, their relevance for life-cycle consumption decisions surely deserves closer investigation.

In this paper, we therefore study the role of consumption externalities in intertemporal consumption choice. We begin by proposing a theoretical model of consumption that allows for peer-group effects. Instead of introducing an ad hoc behavioral model, we extend the standard life-cycle model to account for the notion of "keeping up with the Joneses" in individual felicity functions. Thus, our model is well-grounded in the usual suppositions about human greed and rationality and fully consistent with the forward-looking utility maximization framework as presented, for example, in Browning and Crossley (2001). Apart from providing a coherent economic foundation for our exercise, doing so has the additional advantage that our results are easily compared to those obtained from more traditional versions of the model. Specifically, our specification nests the standard power utility model as well as its demographics-augmented variant as special cases. This allows us to construct simple tests for peer-group effects within the traditional life-cycle framework.

In order to evaluate the model empirically, we derive its first-order condition, an extended version of the well-known consumption Euler equation, and estimate it using US micro data from the Panel Study of Income Dynamics (PSID). Standard Euler equation estimation uncovers substantial predictable consumption co-movement within peer groups, suggesting the

¹A recent example is Luttmer (2004), who also provides further references to this quite sizable literature.

²Already Deaton (1992) suggests some modifications of the standard life-cycle model to take account of household demographic structure or nonseparabilities between consumption and leisure. Attanasio et al. (1999) represents a nice example of such an approach.

presence of possibly strong consumption externalities. However, results obtained under the usual Euler equation framework have to be taken with caution, because estimation is vulnerable to even minor misspecifications. The issue is related to what Manski (1993, 1995) calls the "reflection problem". Basically, to identify and estimate true externalities, we need to discriminate between two competing hypotheses: Is individual consumption growth actually affected by peer-group behavior (endogenous effects) or does it display co-movement within peer groups merely because individuals share similar unobserved characteristics or suffer similar predictable shocks (correlated effects)? Disentangling these two phenomena constitutes the principal challenge when confronting our model with the data. The solution we propose is based on exploiting a social multiplier. Specifically, we adapt Manski's reflection problem to the case of dynamic Euler equations with endogenous regressors. This step allows us to derive further equilibrium conditions implied by our extended model which can be used to re-assess peer effects in intertemporal consumption and provide a more robust test of their relevance. Perhaps surprisingly, once correlated effects are adequately accounted for, our estimation results indicate that there is not much evidence for substantial consumption externalities.

Apart from the aforementioned early contributions that have inspired this paper, our research is also related to a number of different strands in the more recent literature. First, our analysis extends traditional studies of intertemporal consumption profiles based on micro data by allowing for the presence of peer-group effects. Previous research has shown that a "stripped down" power utility version of the model cannot explain key features of the life-cycle profile of consumption. However, some progress has been achieved with the inclusion of demographic preference shifters, as suggested, for example, by Blundell et al. (1994) or Attanasio et al. (1999). The fact that our approach nests this class of models is helpful in that we can assess the importance of peer-groups effects while controlling for the direct impact of relevant (common) demographics. In the same literature, there have also been a few contributions devoted to "internal" habit formation, i.e. persistent effects of an individual's own consumption experience over time. Yet, the empirical evidence for internal habits is mixed at best (see, for instance, Dynan (2000), Guariglia and Rossi (2002), Alessie and Teppa (2002), and Browning and Collado (2004)). Our approach is different from these contributions, because we focus on "external" habit formation: rather than looking at current consumption relative to past consumption for a given individual, we investigate the relationship between an individual's current consumption and that of her peers.

A second related strand of the literature is concerned with intra-period consumption patterns. In fact, most of the studies investigating peer effects in consumption have looked at commodity demand. Building on theoretical work of Gaertner (1974) and Pollak (1976), Alessie and Kapteyn (1991) and Kapteyn et al. (1997), for instance, have shown that peer effects are important for estimating budget share equations. In essence, their work can be understood as referring to the second stage of a two-stage budgeting procedure. Our study, in turn, focuses on the first stage - the intertemporal allocation of consumption - thereby complementing previous research concerned with intra-period demand systems only.

Third, our analysis should be informative regarding the empirical relevance of recent theoretical contributions that have suggested consumption externalities such as "catching" or "keeping up with the Joneses" as potential solutions to empirical puzzles in macroeconomics and finance.³ Indeed, the macroeconomic literature has readily adopted various kinds of internal and external habit specifications for intertemporal consumption, although evidence from microeconomic studies is scarce or absent.

Lastly, we contribute to the literature on identification and estimation of social effects⁴ by adapting Manski's reflection problem framework to a dynamic setting with endogenous regressors. In particular, we demonstrate how the equilibrium conditions derived in Manski (1993, 1995) translate into additional restrictions for dynamic Euler equations that can be exploited to improve estimation and inference.

The remainder of the paper is organized as follows. Section 2 is preparatory and reviews the standard life-cycle model, which we extend, in section 3, to allow for peer effects. In section 4, we present our data and describe the way we construct peer groups. This is followed by a detailed exposition of econometric issues in section 5, where we derive our model specification and discuss identification and inference. In section 6, we turn to our results and their interpretation. Section 7 concludes with some final remarks. Less instructive derivations and econometric technicalities are relegated to the appendix.

2 The Life-cycle Model

In this section we briefly review the main features of the canonical life-cycle model of consumption. This model, which undoubtedly represents the cornerstone in modern literature on consumption, also forms the conceptual basis for our own study.

Consider the intertemporal optimization problem of an infinitely-lived consumer at time t who faces a riskless asset with real after-tax rate of return R_{t+1} . Assume that the consumer has von Neumann-Morgenstern preferences and derives utility from consumption C , with intra-period felicity function u . Assume further that the consumer's rate of time preference is β . We can then write her maximization problem as

$$\max E_t \left[\sum_{j=0}^{\infty} \beta^j u(C_{t+j}) \right] \quad (1)$$

subject to an intertemporal budget constraint; E_t denotes the conditional expectations operator in time t . The first-order condition for this problem is the familiar Euler equation

$$u'(C_t) = \beta E_t [u'(C_{t+1}) R_{t+1}]. \quad (2)$$

³Prominent examples of this line of research include Abel (1990), Gali (1994), Campbell and Cochrane (1999), Ljungqvist and Uhlig (2000) or Binder and Pesaran (2001).

⁴See e.g. Brock and Durlauf (2001) and Manski (2000) for further references.

The left-hand side represents the immediate loss in utility if the consumer marginally increases her asset holdings in t . The right-hand side is the increase in (discounted expected) utility she obtains from the corresponding extra asset payoff in $t + 1$. At an optimum, marginal gains and losses must be exactly equal. Straightforward as it is, this first-order condition of the intertemporal maximization problem is at the core of the life-cycle model throughout all its variants. Intuitively, a rational and farsighted individual aims at smoothing marginal utility throughout her life.

3 Peer-group Effects in the Euler Equation Framework

The standard way to proceed is to parameterize the felicity function $u(\cdot)$, notably by assuming preferences of the constant relative risk aversion (CRRA) type. Estimation is then based on a log-linearized version of (2). In the present paper, we slightly depart from this practice and instead follow an approach first suggested by Attanasio and Browning (1995). Their idea is to start directly from a model for marginal utility $u'(\cdot)$ (or the natural logarithm thereof) that allows for more flexible while still tractable preference specifications. Hence we are able to test for the importance of peer-group effects without relying on an overly restrictive modelling context. As Attanasio and Browning (1995) emphasize, the approach comes at a low cost, since it is still possible to recover the implied utility function by means of integration, if so desired. Likewise, the approach is well-grounded in consumption theory insofar as the empirical model we postulate nests the standard CRRA case with or without additional preference shifters. To illustrate this point, appendix A1 shows how our framework easily accommodates an extended version of the typical CRRA model with peer effects.

With respect to Attanasio and Browning (1995), our central innovation is to add the possibility of consumption externalities into the model. Thus, apart from the key determinants of marginal utility already considered by them, we here allow marginal utility to be also affected by the current consumption level of likely peers. Specifically, we assume that individual marginal utility $u'(\cdot)$ is characterized by

$$\sigma \ln u'_h \left(C_t^h, \left\{ C_t^j | X_t^j = X_t^h \right\}, D_t^h \right) = D_t^h \theta - \ln \left(C_t^h \right) + \gamma ARITM \left[\ln \left(C_t^j \right) | X_t^j = X_t^h \right], \quad (3)$$

where D_t^h represents a vector of basic household characteristics that act as preference shifters, e.g. family size or the number of children. As mentioned before, including such preference shifters is essential for every attempt to take the model to micro data, because consumption is measured at the household level, whereas theory focuses on the individual.

Next, $\ln C_t^h$ represents household h 's log consumption, while $ARITM \left[\ln \left(C_t^j \right) | X_t^j = X_t^h \right]$ denotes the arithmetic mean of the log consumption levels within household h 's peer group, i.e. among households with the same demographic characteristics X_t^h . Intuitively, while marginal utility is assumed to decline in the individual's own consumption level, we posit that it may

also depend on current peer-group consumption,⁵ capturing notions of status concern or jealousy. Thus, our specification allows for intertemporal consumption complementarities between similar households.⁶ The central parameter of interest is γ , which captures the strength of possible peer effects.

Although it may first seem ad hoc, the specification given by (3) is really but a slight extension of the standard power utility model. Accordingly, the basic "stripped down" CRRA model is nested as a special case for $\theta = \gamma = 0$. This is important, since it allows us to construct simple t- or F-tests for consumption externalities against more traditional alternatives.

Combining (3) with a linearized version of the general Euler equation (2), we obtain⁷

$$\Delta \ln \left(C_{t+1}^h \right) = \alpha + \Delta D_{t+1}^h \theta + \sigma \ln R_{t+1} + \gamma ARITM \left[\Delta \ln \left(C_{t+1}^j \right) | X_t^j = X_t^h \right] + \varepsilon_{t+1}, \quad (4)$$

where α contains both the logarithm of the discount rate, β , and higher-order terms stemming from the linearization. The demographic preference shifters D_{t+1}^h now show up in differences, corresponding to the notion that changes in, say, household size or the number of children should have an effect on the household's growth rate of consumption.

It seems worthwhile to provide some intuition for the above Euler equation. After controlling for the effect of demographics, household consumption growth is seen to depend on the interest rate and on average peer-group consumption growth. The effect of the interest rate, on the one hand, reflects an intertemporal substitution motive standard in intertemporal Euler equations. Peer-group consumption growth, on the other hand, is included because households may aim at smoothing their own consumption profile *relative to* that of their peers. One important insight from (4) is that the goal of "keeping up with the Joneses" does not imply excessive current consumption to increase social status. Rather, since the intertemporal budget constraint requires any increase in current consumption to be balanced against lower future consumption, rational forward-looking individuals attempt to maintain their relative position within their peer group as a means of smoothing their marginal utility.

At this point, one additional modification is necessary for us to be able to achieve our goal of testing for peer-group effects. In fact, from looking at (4), one might (and should) be concerned that estimates of γ will pick up spurious correlation rather than true consumption externalities. Specifically, direct effects of the stratification variables X on chosen consumption growth could be falsely interpreted as evidence for peer effects. As an example, one might think that different degrees of education imply different degrees of impatience, i.e. higher or lower discount rates β . As such a phenomenon would concern the whole peer group, similar

⁵In a previous version of the paper, we also analyzed a possible effect of lagged peer-group consumption on an individual's current marginal utility. However, our estimations, which are readily handled within a standard IV framework, did not yield any evidence for such "catching up with the Joneses" in individual preferences.

⁶Note that (3) implicitly treats peer-group consumption at the household level as a potential determinant of marginal utility. Theoretically, it is possible to also incorporate demographic preference shifters in the peer-group term so as to adjust for household size. However, such a model would seem to cause formidable inference problems in practice, while it is not even entirely clear a priori which of the two approaches is more plausible.

⁷See appendix A2 for a complete derivation of equation (4).

behavior could easily be mistaken as evidence for peer effects, whereas the true explanation rests on correlated effects related to observable demographics. In order to distinguish between these two potential phenomena, we control for the direct effects of our stratification variables by including them as additional regressors. This approach neatly accommodates two different strands of the literature. First, we comply with the "reflection problem" framework proposed by Manski (1993, 1995) to allow for both endogenous and correlated effects in what Manski refers to as a "linear endogenous-effects model". Second, we take up the reasoning put forward in part of the consumption literature⁸ that, apart from the preference shifters D_t^h already introduced above, demographics also have to be used to allow a more flexible specification of the discount rate. Thus, by including the additional term $X_t^h \lambda$, we implicitly parameterize $\ln \beta$, which was previously buried in the intercept.

A different issue is the likely presence of common unpredictable shocks, e.g. shocks to the income of particular groups or sudden correlated changes in preferences. At first sight it might seem that such shocks would confound our analysis. However, none of these unpredictable events represent a problem for our purposes, given that peer group consumption - an endogenous variable - will be instrumented with lagged information throughout. In this sense, importantly, we are only dealing with predictable or planned co-movement of consumption.

Our new, augmented Euler equations reads as

$$\Delta \ln \left(C_{t+1}^h \right) = \alpha + \Delta D_{t+1}^h \theta + \sigma \ln R_{t+1} + \gamma ARITM \left[\Delta \ln \left(C_{t+1}^j \right) | X_t^j = X_t^h \right] + X_t^h \lambda + \varepsilon_{t+1}. \quad (5)$$

Equation (5) is the starting point for our estimation strategy. As a practical matter, we will compute cell averages for a given set of discretized stratification variables to obtain nonparametric estimates for $ARITM \left[\cdot | X_t^j = X_t^h \right]$. Note that we must compute these cell averages for each year t . This implies that estimated peer-group means of any endogenous variable have to be treated themselves as endogenous variables with respect to the time dimension. Consequently, such variables will have to be instrumented. This is important to keep in mind when studying intertemporal Euler equations, which typically comprise several endogenous regressors. In addition, replacing the peer-group mean, $ARITM \left[\Delta \ln \left(C_{t+1}^j \right) | X_t^j = X_t^h \right]$, by a first-step estimate, $ARITM_N \left[\Delta \ln \left(C_t^j \right) | X_{t+1}^j = X_t^h \right]$, gives rise to a generated regressor problem.

It is worth emphasizing that the consumers' Euler equations are necessary conditions for any equilibrium within our framework. Building the analysis on such Euler equations, therefore, has the considerable advantage that it allows us to estimate the relevant preference parameters without a full characterization of the particular equilibrium. However, a requirement for identification in this context is that there exist at least one variable that affects individual but not peer-group consumption growth. In our case, the necessary variation is provided by within-peer-group variation in the demographic preference shifters and/or the after-tax interest

⁸Prominent examples include Lawrance (1991), Attanasio and Browning (1995) and Dynan (2000).

rate. Specifically, the demographic preference shifters exploit the fact that consumers aim at smoothing expected discounted marginal utility, while the potential peer effect is formulated with respect to household consumption. After-tax interest rate heterogeneity, in turn, exploits potential differences in the intertemporal substitution motive within peer groups. These two features also distinguish the model in (5) from tests of full consumption insurance with perfect capital markets and pareto-efficient equilibria (see e.g. Mace (1991) and Cochrane (1991)).

At the same time, it is interesting to note that potential peer effects would be identified even in the case of full within-peer-group consumption insurance, perfect capital markets and pareto-optimal equilibrium selection, as long as we analyze externalities based on peer-group consumption levels rather than marginal utility. Certainly, dropping this assumption would result in a loss of identification in a complete-markets setup. Even then, however, identification could be easily restored within a slightly modified framework, e.g. by allowing for within peer-group heterogeneity in the parametric time preference rate.⁹ Before delving deeper into these and other identification and estimation issues, we first introduce our data set and the method used to construct peer groups.

4 The Data

This section briefly describes the data we use for our study. In addition, we provide a comprehensive discussion of how we define and construct peer groups. We also report some descriptive statistics about the final sample used in the estimation.¹⁰

4.1 The Panel Study of Income Dynamics (PSID)

Our analysis is based on data from the well-known Panel Study of Income Dynamics (PSID). The PSID is a longitudinal survey of a representative sample of US individuals and their families. It has been ongoing since 1968. The data were collected annually through to 1997, and biennially starting in 1999. As a consequence of low attrition rates, the success in following young adults as they form their own families, and recontact efforts, the sample size has grown from 4,800 families in 1968 to more than 7,000 families in 2001. While the PSID has a very broad content, including economic and demographic as well as sociological and psychological measures, its coverage of consumption behavior is relatively limited. Indeed, the only measure of consumption available in the files is food consumption (at home and in restaurants), which, moreover, was a recurrent item in the survey questionnaire only between 1974 and 1987. Accordingly, the PSID offers a maximum of 14 consecutive annual observations to investigate households' intertemporal consumption patterns.

⁹Such a strategy would work since full consumption insurance only implies that the discounted growth rate of marginal utility is constant across households. Thus, heterogeneity in discount rates within peer groups is sufficient to generate time-varying paths for marginal utility, which provide the necessary variation to identify the model.

¹⁰Details related to necessary data cleaning procedures can be found in appendix A3.

Having data only on food consumption is an obvious drawback. Specifically, for our exercise to be valid, we need to assume separability of utility between food consumption and other expenditure items. However, we have to make do with the data available, and since the use of actual panel data covering a reasonably long horizon is crucial for our analysis, we see no alternative to using the PSID. In particular, synthetic panel data, which are also frequently used to estimate consumption Euler equations, are not an option. The reason is that such data are already based on cohort aggregation, thus eliminating one of the key dimensions along which consumption externalities should be analyzed. Similarly, imputation techniques like those suggested by Skinner (1987) and Blundell et al. (2005) do not provide a superior alternative in our case, since identification of possible peer effects would ultimately rely on rather artificial individual-level variation among members of the same socioeconomic group. These concerns basically preclude the use of the Consumer Expenditure Survey (CEX), with its very short four-quarter panel dimension, as a potential data source. Assuming separability and focusing on food expenditure therefore seems to us a necessary price to be paid, at least as long as long-horizon panel data containing broader consumption measures are not available.¹¹

As mentioned, the data we use cover the interview period 1974 to 1987. After necessary deletions, including all households in the so-called "poverty subsample", and some data losses due to implausible or missing observations, all of which are duly documented in appendix A3, our data set still comprises roughly 26,000 observations. Based on this sample, we next start constructing peer groups.

4.2 Peer-group Construction

The specification of peer groups is critical for any analysis of social interactions. Most importantly, the reference groups we choose to consider are taken to be characteristic of an individual's social environment, so we must take a stance on what personal attributes plausibly define such an environment. The ideal solution would be to use observed behavior and infer the most relevant determinants or dimensions of social reference groups directly from the data. However, as pointed out by Manski (1993, 1995) in his seminal contribution on endogenous social effects, such an approach would render identification impossible and make the social effects model hold tautologically. Thus he concludes that "informed specification of reference groups is a necessary prelude to analysis of social effects" (Manski (1993), p. 536). Naturally, parametric identification could also be attained through specific functional form restrictions on either the model or the way reference group characteristics are aggregated. We do not follow this approach here, since in our application functional form restrictions would have neither a compelling basis in theory nor an intuitive interpretation. Instead, we borrow results from the literature on group processes and social comparison in social psychology to motivate our sample stratification. Studies of social comparison processes (see, for example, Festinger (1954))

¹¹The lack of better panel data has led to a substantial body of literature using food expenditure to explore consumer behavior. Examples include Hall and Mishkin (1982), Zeldes (1989), Lawrance (1991), Runkle (1991), Jacobs (1999) or Dynan (2000), to name just a few.

emphasize that people primarily compare themselves to members of their own social group, i.e. to individuals who are similar along dimensions such as age, gender or education. As in Kapteyn et al. (1997), we will therefore treat individuals that share such basic characteristics as relevant reference groups.

Given the focus of our exercise, it seems important to account for characteristics related to social achievement and status. These clearly include age and education - two categories used in a study by Woittiez and Kapteyn (1998) - but we prefer to also consider other characteristics that are relevant to an individual's self-conception like race, gender or urbanity. Hence, we construct reference groups based on the following attributes characterizing household heads in our sample: age cohort, race, gender, educational attainment, and "size of the nearest city" as a measure for urbanity.¹² Specifically, cell averages are computed using six-year cohorts based on the household head's age in 1974; a dummy indicating whether the household head is white or non-white; a gender dummy; a categorical education variable that takes on one of three different values depending on whether the household head has had less than high school, a high school degree or a completed college education; and lastly a city size variable that indicates whether the nearest city has less than 50,000, between 50,000 and 500,000, or more than 500,000 inhabitants. Obviously, the list of strata-defining characteristics could easily be extended. In a sample of limited size, however, this has to be traded off against the disadvantage of ending up with overly small reference groups or substantial data losses. We therefore settle for the above-mentioned five criteria.

In order to obtain a meaningful proxy for peer-group means, we consider only strata consisting of at least 15 households in a given year. This choice again represents a compromise between different goals. While larger cell sizes are in principle desirable, they would also imply more data losses for a given set of stratification variables, so we have to strike a balance.¹³ Even so, the need to delete observations pertaining to overly small reference groups causes a reduction of our sample to a final size of 18,126 observations. Table 1 indicates the number of observations per year contained in the final sample. Similarly, table 2 provides information on how many peer groups remain in the sample each year and what minimum, maximum, average and median size they have. In addition, tables 3 and 4 contain basic summary statistics for our sample before and after the deletion of households belonging to small peer groups. Comparison of these tables shows that our sample becomes "more white", "more male" and somewhat "more educated" and "less urban" as a consequence of data deletions, whereas compositional

¹²These same variables have also been found to be important predictors for individual welfare functions over income as studied in the Leyden approach (see e.g. Van Praag and Frijters (1999)), which provides additional support for our stratification strategy.

¹³Note, however, that our exercise is not subject to the minimum cell size requirements typical of studies based on synthetic cohort techniques. When constructing synthetic cohorts, researchers inevitably need large cell sizes to minimize sampling variability in the proxy for consumption growth. Here, we do not face this problem as we use true panel data. Our only concern is to avoid peer-group means being overly affected by a few households that could be outliers. For this purpose, using a threshold of 15 seems acceptable. Importantly, our estimation explicitly takes account of sampling variability at the peer-group level. Moreover, we conduct robustness checks to assess the sensitivity of our results to the minimum group size chosen.

changes in terms of birth cohorts are minor. Essentially, imposing a lower bound on cell size removes most of the households headed by non-whites and women. Although this is of course unfortunate, we still find our final sample relatively well-balanced even with respect to other studies that do not face data constraints associated with peer-group construction. For example,

Table 1: Number of observations by year in final data set*

Year	Number of Observations
74-75	1,278
75-76	1,357
76-77	1,379
77-78	1,388
78-79	1,312
79-80	1,331
80-81	1,413
81-82	1,389
82-83	1,372
83-84	1,414
84-85	1,412
85-86	1,469
86-87	1,612
Total	18,126

*Baseline case: Minimum group size: 15.

Table 2: Number and size of reference groups by year

Year	No. of Groups	Min. Size	Max. Size	Mean Size	Median Size
74-75	46	15	73	27.8	24
75-76	46	15	82	29.5	26
76-77	48	15	83	28.7	24
77-78	47	15	85	29.5	24
78-79	47	15	89	27.9	24
79-80	44	15	102	30.3	24
80-81	46	15	104	30.7	24
81-82	42	15	109	33.1	25
82-83	40	16	112	34.3	27
83-84	36	15	115	39.3	33
84-85	36	15	112	39.2	33
85-86	36	16	115	40.8	31
86-87	39	15	117	41.3	31

Table 3: Summary statistics for data set before deletion of small groups (in %)

Race	white	90.56
	non-white	9.44
Gender	female	13.89
	male	86.11
Education	less than high school	21.79
	high school or more	36.33
	finished college or more	41.88
City Size	less than 50,000	39.55
	between 50,000 and 500,000	38.08
	more than 500,000	22.37
Cohort (age in 1974)	5-10	0.03
	11-16	2.54
	17-22	15.91
	23-28	23.20
	29-34	13.59
	35-40	10.11
	41-46	13.02
	47-52	11.50
	53-58	7.35
	59-65	2.75
Total Number of Observations		26,358

Table 4: Summary statistics for final data set (in %)

Race	white	99.91
	non-white	0.09
Gender	female	0.42
	male	99.58
Education	less than high school	14.67
	high school or more	36.89
	finished college or more	48.44
City Size	less than 50,000	45.44
	between 50,000 and 500,000	38.09
	more than 500,000	16.46
Cohort (age in 1974)	5-10	0.00
	11-16	1.71
	17-22	17.00
	23-28	26.24
	29-34	14.98
	35-40	9.10
	41-46	13.50
	47-52	12.19
	53-58	4.18
	59-65	1.10
Total Number of Observations		18,126

many authors in the consumption literature have dropped female-headed households from their sample straight away, thus obtaining the same selectivity we are faced with as a result of our necessary data deletions. In conclusion, our focus on households belonging to sufficiently big reference groups implies a natural qualification on the interpretation of our results, insofar as we cannot extrapolate to other subpopulations. Nevertheless, our sample provides interesting insights about sizeable and important strata of society.

As was stressed before, it is inherent to the nature of the reflection problem that specific stratification approaches cannot be informed or judged by observed behavior without leading to tautology. It may still be informative to take a first look at the correlation between individual consumption growth and the reference-group counterpart that we have constructed. In fact, the raw partial correlation is 0.2, more than three times as high as the correlation between individual and aggregate annual consumption growth in the whole sample (0.06).

5 Specification, Identification, and Inference

5.1 The Reflection Problem and Omitted Correlated Effects

In principle it is possible to estimate an equation like (5) using instrumental variables (IV). Thus we will report the corresponding results below. However, estimation is a rather delicate issue in this case. In particular, results are probably very sensitive to even minor misspecification (or omission) of direct demographic effects. To be sure, we can (and do) control for direct effects from our stratification variables by including these dummies as additional regressors. This corresponds to a parameterization of the conditional expectation of the error term as suggested by Manski (1993) and can be interpreted as an auxiliary parametric model for the consumer's discount rate. Yet, such direct controls are necessarily imperfect. For one thing, effects could stem from complicated interactions of the demographic information we use. Further, consumption growth might be affected by additional demographic factors omitted from the model. This will remain a potential problem, even if a specification already takes into account all variables that have been identified as relevant in the previous literature. In essence, there is no safe guidance as to what precise set of demographic variables have to be included in taking the life-cycle theory to the data. In most applications, this point may be a purely academic one. For our study, however, it is critical, because it further raises the challenge of discriminating between true consumption externalities and merely correlated effects.

Basically, any omission of direct demographic effects in (5) is likely to cause an upward bias in the estimate for γ , combined with an uninformative J-statistic. The reason lies in the mechanics of IV estimators. Recall that estimates are obtained from minimizing the (weighted) correlations between instruments and residuals, i.e. *estimated* errors. The estimator tends to purge such components from the residual that are correlated with the instruments. In the case of equation (5), this may imply that omitted demographic effects or predictable "common

shocks” within peer groups are spuriously eliminated by assigning a value near 1 to γ . The social interactions term thus picks up any omitted effects present at the level of peer groups. Moreover, although the model is clearly misspecified in this case, the misspecification would be virtually impossible to detect. In order to understand why the test of overidentifying restrictions may fail, note that it is based on the minimized objective function of the estimator. Its power to reject a given specification prevails only to the extent that the estimated error term actually displays sufficient correlation with the instruments. Given the above reasoning, it seems fair to suspect that the J-test may have very low power.

As a bottom line, specification (5) is very vulnerable to even minor omissions of relevant demographic information. Above and beyond what this would imply for any analysis of micro consumption data, it poses the very concrete problem here that our main coefficient of interest might easily be upward biased, thus jeopardizing any conclusions about peer-group vs. correlated effects. Importantly, this is true despite the fact that the model is well identified under the assumption of correct specification and the availability of valid instruments.¹⁴

However, the situation is not quite as unfortunate as it might first seem. The solution we propose relies on the fact that optimal consumption growth rates need to be consistent within peer groups. This insight provides us with a set of additional equilibrium restrictions that can be exploited to discriminate between our two hypotheses of interest. Specifically, the additional equilibrium conditions allow us to transform our model and obtain a new specification which does not suffer from the aforementioned shortcomings. The general idea of exploiting an internal consistency argument goes back to Manski’s (1993, 1995) contributions on how to circumvent the reflection problem in the identification of endogenous social effects. Here we adapt Manski’s framework to the case of IV estimation and inference for Euler equations featuring contemporaneous consumption externalities.

First, by aggregating equation (5) for each peer group separately, we obtain

$$\begin{aligned}
 ARITM \left[\Delta \ln \left(C_{t+1}^j \right) | X_t^j = X_t^h \right] &= \alpha + ARITM \left[\Delta D_{t+1}^h | X_t^j = X_t^h \right] \theta & (6) \\
 &+ \sigma ARITM \left[\ln R_{t+1} | X_t^j = X_t^h \right] \\
 &+ \gamma ARITM \left[\Delta \ln \left(C_{t+1}^j \right) | X_t^j = X_t^h \right] \\
 &+ X_t^h \lambda + ARITM \left[\varepsilon_{t+1} | X_t^j = X_t^h \right],
 \end{aligned}$$

which is referred to as a ”social equilibrium condition” for each stratum. This additional equilibrium restriction recognizes the fact that the consumption growth terms on both sides of the Euler equation have to be mutually consistent.

¹⁴Note, in fact, that our exercise does not suffer from a fundamental identification problem like the one described in Manski (1993, 1995). In his model, all explanatory variables affect the outcome variable directly as well as through their respective reference group levels. The consequence is perfect collinearity between the peer-group averages and the other explanatory variables rendering identification impossible. In our case, as we can safely rule out what Manski coins ”contextual effects”, we obtain an exclusion restriction that allows us to identify the model.

Next, we assume that γ is not exactly equal to 1. Rearranging terms allows us to write

$$\begin{aligned}
ARITM \left[\Delta \ln \left(C_{t+1}^j \right) | X_t^j = X_t^h \right] &= \frac{1}{1-\gamma} \alpha & (7) \\
&+ \frac{1}{1-\gamma} ARITM \left[\Delta D_{t+1}^h | X_t^j = X_t^h \right] \theta \\
&+ \frac{\sigma}{(1-\gamma)} ARITM \left[\ln R_{t+1} | X_t^j = X_t^h \right] \\
&+ \frac{1}{1-\gamma} X_t^h \lambda + \frac{1}{1-\gamma} ARITM \left[\varepsilon_{t+1} | X_t^j = X_t^h \right].
\end{aligned}$$

Thus, the social equilibrium condition (7) implies that, for each population stratum, peer-group consumption growth depends only on the peer-group means of the other explanatory variables, i.e. averages of family size changes, the average log after-tax interest rate and demographics dummies. In order to exploit these additional restrictions, we combine condition (7) with the augmented Euler equation (5) to obtain

$$\begin{aligned}
\Delta \ln \left(C_{t+1}^h \right) &= \alpha + \Delta D_{t+1}^h \theta + \frac{\gamma}{1-\gamma} ARITM \left[\Delta D_{t+1}^h | X_t^j = X_t^h \right] \theta & (8) \\
&+ \sigma \ln R_{t+1} + \frac{\gamma \sigma}{(1-\gamma)} ARITM \left[\ln R_{t+1} | X_t^j = X_t^h \right] \\
&+ \left(1 + \frac{\gamma}{1-\gamma} \right) X_t^h \lambda + u_{t+1},
\end{aligned}$$

where u_{t+1} has been introduced as a shorthand notation for the combined error term $\varepsilon_{t+1} + \frac{1}{1-\gamma} ARITM \left[\varepsilon_{t+1} | X_t^j = X_t^h \right]$. Equation (8) will serve as our principal estimating equation. Of course, in the actual estimation we replace the peer-group means $ARITM \left[\cdot | X_t^j = X_t^h \right]$ with the corresponding nonparametric estimates $ARITM_N \left[\cdot | X_t^j = X_t^h \right]$ based on our sample stratification.

Note that (8) provides a much improved basis to estimate actual peer effects and properly assess model specification. Above all, the right-hand side of the equation no longer includes endogenous peer-group consumption growth as a regressor. Recall that this term is the source of concern in our initial equation (5), since we suspect that it spuriously picks up any predictable group-specific components from the error term.

Some intuition should be provided regarding the way consumption externalities operate in (8). In fact, these externalities are now estimated from a social multiplier, i.e. the indirect effects on a peer's optimal consumption growth operating through peer-group averages of the standard explanatory variables. To give an example, interest rates are one theoretically undisputed determinant of consumption growth. To the extent that higher average interest rates raise average consumption in an individual's peer group, they also cause the individual herself to raise consumption if peer-group effects are present. This explains why the coefficient in front of the peer-group interest rate variable contains γ . At the same time, equation (8) naturally accounts for the direct effects of the standard explanatory variables. Specifically, they are iden-

tified provided that there is some household-level variation relative to the peer-group averages. Consider again the example of interest rates. Many authors have estimated consumption Euler equations using pre-tax interest rates. Apart from the general qualms one may have about this approach, following it would actually make it impossible for us to distinguish between the effect on consumption of the interest rate faced by the individual herself and the one faced by her peers. Both interest rate terms in (8) would be identical. This gives us a strong rationale for using after-tax interest rates R_{t+1} as regressors. Finally, we can also identify correlated effects as captured in X_t^h . The coefficient γ being identified from the social multipliers, we can isolate the direct impact of the stratification variables by netting out, from the total effects $\left(1 + \frac{\gamma}{1-\gamma}\right) \lambda$, any indirect effects $\left(\frac{\gamma}{1-\gamma}\right) \lambda$ that stem from the consumption externality. Thus, the reflection problem framework allows us to conduct proper inference with respect to all parameters of interest.

As a practical matter, however, the above discussion also suggests that estimation remains a challenge, especially since the variation we can exploit to estimate consumption externalities must come from within peer groups. Moreover, some of the explanatory variables such as the interest rate are endogenous. These variables (and their respective peer-group averages) need to be instrumented, which reduces the necessary variation even further. Hence, it may be difficult to obtain very precise estimates of the parameters of interest.

5.2 Explanatory Variables and Instrument Choice

5.2.1 Explanatory Variables

In order to estimate Euler equations on micro data, it is essential to properly account for real-life heterogeneity and demographic variation. Indeed, the basic life-cycle model is formulated at the level of an ahistorical individual, whereas actual data refer to households with specific demographic patterns and a finite lifetime. At a minimum, most economists have considered changes in household size as an important determinant of consumption growth.¹⁵ We follow the literature by including family size, the number of major adults and the number of children as preference shifters D_{t+1} in our parameterization for marginal utility, (3). Accordingly, the Euler equations (4), (5) and (8) include changes in these demographic variables as additional regressors.

Moreover, as already discussed above, we take up the reasoning of Lawrance (1991) and others and include further demographics directly in the Euler equation to allow for differences in time preference rates across different subpopulations. Note that this is crucial in our application in order to distinguish endogenous social effects from correlated effects, i.e. direct effects of the stratification variables on the dependent variable. Specifically, since Lawrance (1991) argues that some of our stratification variables could be associated with large differences in time preference rates across subpopulations, we include all our stratification variables as explanatory

¹⁵Prominent examples include Attanasio and Weber (1993, 1995), Attanasio and Browning (1995), Attanasio et al. (1999) or Dynan (2000). See Deaton (1992) or Attanasio (1999) for an overview.

variables in equation (5). Hence, X_t comprises race, gender, education and cohort dummies as well as the city size dummies for different degrees of urbanity. As a proxy for the riskless asset return R_{t+1} , we use the real after-tax one year US T-Bill rate, constructed as the average of twelve year-to-year rates. Lastly, we construct peer-group averages of all the explanatory variables since these are needed for estimating (8).

5.2.2 Instrument Choice

All of our estimating equations contain some endogenous regressors. In particular, the real after-tax interest rate in (5) and (8) is an endogenous variable that needs to be instrumented. Furthermore, the peer-group mean of log consumption growth in equation (5) naturally needs to be instrumented, as well. Within a forward-looking, rational expectations framework like the one considered here, every variable that is contained in the current information set basically provides a valid instrument. A qualifier is necessary insofar as measurement error in levels or time-aggregation can lead to autocorrelation in growth rates, thus invalidating the first lag of, say, income growth as an instrument for current income growth. Apart from these considerations on instrument validity, we try throughout to pick instruments that are likely to contain a lot of information about the endogenous variables. The goal is to have high predictive power for our endogenous regressors without excessive instrumentation, i.e. without recourse to many (weak) instruments that simply drive up the degrees of freedom. Moreover, we attempt to attain a reasonable balance between aggregate and individual-specific variables. Thus, apart from all exogenous variables, we use four lags of the real T-Bill rate, four lags of real stock returns from the S&P 500, the second to fourth lag of the CPI inflation rate, the second and third lag of real income growth and their squares and cubes as well as lagged labor market status of head and spouse as instruments.

6 Estimation Results

6.1 Results for Equation (5)

We start by estimating (5) using two-step GMM. The exact formulation of the estimator is described in appendix A4. This appendix also details how we account for sampling variability in the estimated peer-group means, a variant of the "generated regressor" problem discussed, for example, in Newey and McFadden (1994).

Table 5 presents point estimates along with appropriately adjusted standard errors in parentheses. Note first that the estimates tend to lie in a reasonable range, with confidence intervals sufficiently small to pin down parameter values quite precisely. The coefficients pertaining to changes in household size are all of the expected sign and size. For instance, consumption growth is estimated to rise by roughly 20 %, *ceteris paribus*, with the arrival of a new major adult in the household, whereas one additional child increases consumption growth by less than 10 %. The intertemporal elasticity of substitution σ appears relatively small, taking a value

of 0.06. Most importantly, however, the parameter associated with peer-group consumption growth, γ , indicates strong consumption externalities, with a point estimate of 0.96 and a small standard error. In addition, the usual specification tests lend support to these results in that the J-statistic clearly fails to reject the model at any conventional level of significance.

Table 5: Euler equation estimates incl. peer-group consumption growth

Parameter	
Δ family size	0.091 (0.0063)
Δ major adults	0.103 (0.0172)
Δ children	-0.018 (0.0073)
σ	0.061 (0.0139)
γ	0.955 (0.0270)
J-statistic	9.23
p-value	0.7556

Estimates account for the presence of generated regressors. Standard errors are in parentheses. The estimation also includes an intercept as well as direct controls for the stratification variables.

Basically, the results in table 5 indicate that the instruments we use are orthogonal to deviations of individual consumption growth from its reference group mean, controlling for other typical regressors. In other words, while consumption within subgroups seems to show substantial predictable co-movement, deviations from peer-group consumption growth appear largely unpredictable. Although we have already argued that it is impossible to infer reference group characteristics from observed behavior, this preliminary result is still noteworthy. It suggests that there are important predictable trends at the level of the groups we have chosen to consider. Whether or not the cause lies in actual peer effects, however, has to be investigated using a framework that is more robust to potential misspecification. Indeed, it seems doubtful whether we should take the high estimates of γ at face value. Given the structure of equation (5), if our specification of demographic controls is incomplete, omitted correlated effects may easily be mistaken for true consumption externalities. At the same time, the problem may not become apparent from standard J-tests. Fortunately, the above social equilibrium conditions provide additional restrictions that we can exploit to obtain more reliable estimates for γ and the other parameters of the model.

6.2 Results for Equation (8)

Therefore we next turn to the estimation of (8). The first column of table 6 displays the results for our baseline specification, which again includes a full set of demographic controls in order to capture correlated effects associated with our stratification variables. All of the estimates in table 6 account for the presence of generated regressors.

Table 6: Euler equation estimates imposing the social equilibrium condition

Parameter	Baseline Specification	Excess Sensitivity Test
Δ family size	0.096 (0.0065)	0.095 (0.0066)
Δ major adults	0.103 (0.0183)	0.106 (0.0192)
Δ children	-0.019 (0.0075)	-0.019 (0.0075)
σ	0.150 (0.0691)	0.153 (0.0702)
γ	0.111 (0.1803)	0.101 (0.1841)
log income growth		-0.012 (0.0264)
J-statistic	46.82	46.18
p-value	0.0001	0.0000

All estimates account for the presence of generated regressors. Standard errors are in parentheses. The estimations also include an intercept as well as direct controls for the stratification variables.

Note first that the estimated effects of changes in family composition are virtually identical to the ones estimated from (5) before. However, imposing the social equilibrium conditions alters the results with respect to interest rate effects and, most strikingly, the consumption externality. The direct interest rate effect is now estimated to be somewhat larger with a point estimate of 0.15, more in line with existing results from the literature. Even more interestingly, peer-group effects appear to be small and insignificant now, the parameter estimate being just 0.11 with a standard error of 0.18. Hence the 95 % confidence interval clearly excludes values of γ above 0.5, let alone such close to one. This is nothing short of a reversal of the results suggested by table 5 above: once we conduct inference in a framework that is more robust to misspecification, the empirical support for peer effects vanishes almost completely. In addition, the J-test now rejects the model at the 1 % level, indicating that there are predictable changes in log consumption growth which cannot be fully explained by the regressors contained in (8). Thus, the overidentifying restrictions suggest evidence for omitted variables or, put differently, against the rational expectations life-cycle model in its current specification.

Before addressing the implications of this finding in greater detail, it seems informative to

also subject our model to an "economic" test. The second column of table 6 reports results from a specification that includes instrumented current money income growth as an additional regressor. This specification is often referred to as an "excess sensitivity test" of the life-cycle model. In fact, the theory of intertemporal optimization implies a coefficient of zero for the added regressor, indicating no impact of predictable income changes on consumption growth. As can be seen from the table, the coefficient we estimate is indeed small and insignificant, while the estimates for all other parameters remain virtually unchanged. Hence our model specification passes this "economic" test of the life-cycle model: there is no evidence for consumption changes that are related to predictable changes in income.

Lastly, table 7 checks the robustness of our results with respect to different minimum group sizes. We re-estimate our baseline model based on data with minimum cell sizes of 10, 20 and 25, respectively. The resulting point estimates are nearly the same as before, with standard errors increasing as we increase the minimum cell size, most probably because this implies a loss of observations. Nevertheless, all of our previous findings are solidly confirmed.

Table 7: Robustness checks for baseline specification

Parameter	Minimum Cell Size: 10	Minimum Cell Size: 20	Minimum Cell Size: 25
Δ family size	0.097 (0.0063)	0.098 (0.0070)	0.104 (0.0085)
Δ major adults	0.096 (0.0178)	0.102 (0.0187)	0.099 (0.0202)
Δ children	-0.023 (0.0073)	-0.023 (0.0081)	-0.029 (0.0098)
σ	0.172 (0.0707)	0.124 (0.0677)	0.122 (0.0949)
γ	0.061 (0.1861)	0.163 (0.1801)	-0.054 (0.3330)
J-statistic	55.82	38.71	28.27
p-value	0.0000	0.0012	0.0294
Number of Observations	19,843	16,179	13,333

All estimates account for the presence of generated regressors. Standard errors are in parentheses. The estimations also include an intercept as well as direct controls for the stratification variables.

6.3 Discussion

Taken together, our results suggest that the prima facie evidence in favor of peer effects in intertemporal consumption has to be taken with great caution. Certainly, our initial estimates from equation (5) unveiled strong predictable co-movement of consumption within peer groups.

This can arise either from actual complementarities in consumption or from omitted factors that are not captured by the theoretical specification we consider. The more trustworthy estimates from equation (8) clearly confirm the latter explanation: once we exploit Manski's reflection problem framework to improve inference, the empirical support for peer effects appears weak at best. Specifically, the estimated values for γ are small and insignificant throughout.

In this context, the statistical rejection of our baseline model equation is actually quite instructive in that it helps to explain the apparent contradiction between the results from our initial regressions (table 5) and those from the estimation of the transformed model (table 6). They are, in fact, two sides of the same coin: If the estimated peer-group effects in table 5 are really due to omitted factors, the model is misspecified, even though the J-test may fail to show it. In this situation, estimation of the transformed model (8) should not only reveal the omitted variable bias (through lower estimates of γ) but also improve the power of the associated specification tests. This is precisely what we find.¹⁶

The result of no or mild peer effects may appear surprising. One should note, however, that our analysis is focused on the precise theoretical implications of the life-cycle model. Specifically, the model identifies peer effects with a tendency of smoothing consumption profiles relative to those of relevant peers. This implication, which we cannot confirm in the data, may be quite distinct from what, in a casual discussion, would be associated with the notion of peer effects. Thus, the intuitive plausibility of peer effects may not be all that hard to square with our specific negative finding. For one thing, consumption externalities may still be important for the intratemporal allocation of expenditure on specific goods, as indicated by the evidence in Kapteyn et al. (1997). In addition, it is conceivable that peer effects would be confirmed in behavioral models that depart from a forward-looking rational expectations framework, e.g. by allowing for myopic "overspending". The latter point is important, because our analysis leaves open to what extent the life-cycle framework itself may be too restrictive. While J-tests of our preferred equation (8) imply a statistical rejection of the model, it passes the economic excess sensitivity test.

The critical touchstone is the presence of predictable components in consumption growth that are not captured by any of the explanatory variables. However, it is unclear whether or not the life-cycle model could be further extended, in a convincing fashion, to account for these omitted factors, or whether there is a more fundamental flaw in the framework. One view is that the core predictions of the intertemporal optimization framework are borne out in the data, even if a rich set of preference shifters may have to be added to capture real-world heterogeneity. In this case, our results would simply indicate that even accounting for

¹⁶It should be noted that other economists have estimated Euler equations based on food consumption data from the PSID without rejecting the model. Examples include Zeldes (1989), Runkle (1991) and Dynan (2000). Although our results imply a statistical rejection of the model, we do not deem this finding too surprising. Indeed, the power of the typical Sargan test is very sensitive to both sample size and instrument choice. Specifically, more extensive use of (weak) instruments would clearly work toward a non-rejection of the model in our case, without changing any of our key results. Irrespectively, since our model nests most of the previous estimation approaches in the literature, results remain comparable.

some group-specific demographics already drives down the seeming evidence in favor of peer effects. A different view holds that the life-cycle model keeps missing some essential driving force of households' consumption dynamics and that adding more and more demographics terms would merely immunize the theory against refutation. In line with this view, Souleles (2005) presents evidence that consumer confidence indices have considerable predictive power for consumption growth rates, above and beyond the determinants posited by the standard model. Interestingly, our findings are compatible with this sceptical view, as well. Indeed, the information contained in consumer confidence indices might be seen as one of the sources generating the observed co-movement in consumption within peer groups. Clearly, if such objections to the life-cycle model have merit, one should seek a superior modelling context in which to analyze intertemporal consumption. For lack of obvious alternatives, however, we confine our analysis here to the standard life-cycle setup, which continues to be a building block of modern economics. Its prominence is witnessed not least by several recent contributions that have suggested models of intertemporal optimization with consumption externalities to account for important empirical puzzles in macroeconomics and finance.

Thus, future work may also confront our results with those obtained from alternative identification approaches within the same framework, or simply from new and better data as they become available. Either way it will be important to carry on examining the microeconomic evidence for a modelling device that has become almost commonplace in many macroeconomic applications. Our evidence, in fact, cautions against the view that the relevance of peer group effects in life-cycle consumption can be taken for granted.

7 Conclusion

In this paper we make two contributions. First, we derive a framework for analyzing the importance of peer-group effects, or "keeping up with the Joneses", for the dynamic pattern of household consumption. Second, we confront our model with micro data from the PSID to obtain empirical evidence on such consumption externalities. Our approach has the advantage that it fully nests more traditional versions of the life-cycle model. Specifically, we obtain an otherwise standard Euler equation that allows individual utility to be also affected by peer-group consumption.

Starting with a simple IV-type regression of household consumption growth on average peer-group growth and several controls, we find strong evidence for expected consumption co-movement within reference groups constructed on the basis of age, education, race, gender and urbanity. This suggests some scope for peer-group effects. However, inference at this step is very vulnerable to even minor misspecification, especially the omission of demographic control variables. We argue that additional restrictions are required to adequately discriminate between true consumption externalities and merely correlated effects. Our solution builds on the "reflection problem" framework developed by Manski (1993, 1995), which we adapt for the case of dynamic Euler equations with endogenous regressors. This transformation solves

our inference problem and allows us to distinguish more robustly between the two competing explanations for the observed co-movement of consumption.

Once the possibility of omitted correlated effects is properly accounted for, our estimation results indicate that there is hardly any evidence for peer-group effects in households' intertemporal consumption decisions. It thus appears that standard IV-type estimation based on the extended Euler equation mainly picks up spurious correlation due to neglected common factors. Hence, while our results do not support the hypothesis of strong peer effects within the standard life-cycle framework, we also conclude that this framework requires very flexible specifications to capture the dynamic pattern of households' consumption allocations and make their driving forces visible.

Appendix

A1: Marginal Utility for Isoelastic Preferences with Peer Effects

The following exposition shows how peer effects can be neatly introduced into the standard CRRA utility framework that features prominently in much of the consumption literature. Specifically, we derive an expression for marginal utility which is nested in the general specification we propose in this paper.

Consider the CRRA utility function incorporating peer effects given by

$$u(C_t) = \frac{1}{1 - \frac{1}{\sigma}} \left(\frac{C_t}{\left(GEOM [C_t^j | X_t^j = X_t^h] \right)^\gamma} \right)^{1 - \frac{1}{\sigma}}, \quad (9)$$

where $GEOM (C_t^j | X_t^j = X_t^h)$ denotes the geometric mean of the current consumption levels of an individual's peers. This utility function nests the standard CRRA case for $\gamma = 0$. Note further that we use the geometric mean primarily for the sake of analytical convenience in deriving a linearized Euler equation. However, it could also be motivated by considering its favorable properties with respect to the susceptibility to outliers.

The corresponding marginal utility is therefore

$$u'(C_t) = \left(\frac{C_t}{\left(GEOM [C_t^j | X_t^j = X_t^h] \right)^\gamma} \right)^{-\frac{1}{\sigma}}. \quad (10)$$

Taking logs, we obtain

$$\ln(u'(C_t)) = -\frac{1}{\sigma} \left(\ln(C_t) - \gamma \ln \left(GEOM [C_t^j | X_t^j = X_t^h] \right) \right), \quad (11)$$

or, using the fact that the natural log of the geometric mean equals the arithmetic mean of the logged components,

$$\sigma \ln(u'(C_t)) = -\ln(C_t) + \gamma \left(ARITHM \left[\ln(C_t^j) | X_t^j = X_t^h \right] \right). \quad (12)$$

This representation corresponds with the equation, (3), we use for modelling marginal utility.

A2: Derivation of the Euler Equation

This section follows the derivations in Attanasio and Browning (1995). We start with the general Euler equation (2) given by

$$u'(C_t) = \beta E_t [u'(C_{t+1}) R_{t+1}]. \quad (13)$$

Suppressing the conditional expectations operator, we can rewrite the above Euler equation as

$$u'(C_t) = (\beta u'(C_{t+1}) R_{t+1}) (1 + \tilde{\varepsilon}_{t+1}), \quad (14)$$

where $\tilde{\varepsilon}_{t+1}$ denotes an expectational error with $E_t [\tilde{\varepsilon}_{t+1}] = 0$. Taking logs, we obtain

$$\ln u'(C_t) = \ln \beta + \ln u'(C_{t+1}) + \ln R_{t+1} + \ln (1 + \tilde{\varepsilon}_{t+1}). \quad (15)$$

Note that, by Jensen's inequality, the error term now has nonzero expectation:

$$E_t [\ln (1 + \tilde{\varepsilon}_{t+1})] \leq \ln (E_t [1 + \tilde{\varepsilon}_{t+1}]) = 0. \quad (16)$$

This problem can be dealt with by using a second-order Taylor approximation of $\ln (1 + \tilde{\varepsilon}_t)$ to obtain

$$\begin{aligned} \ln u'(C_t) &= \left(\ln \beta - \frac{1}{2} s^2 \right) + \ln u'(C_{t+1}) + \ln R_{t+1} + \left(\tilde{\varepsilon}_{t+1} + \frac{1}{2} s^2 - \frac{1}{2} (\tilde{\varepsilon}_{t+1})^2 \right) \\ &= \left(\ln \beta - \frac{1}{2} s^2 \right) + \ln u'(C_{t+1}) + \ln R_{t+1} + \zeta_{t+1} \end{aligned} \quad (17)$$

with $s^2 = E_t [(\tilde{\varepsilon}_{t+1})^2]$ for all t and ζ_{t+1} representing a combined error term with $E_t [\zeta_{t+1}] = E_t \left[\left(\tilde{\varepsilon}_{t+1} + \frac{1}{2} s^2 - \frac{1}{2} (\tilde{\varepsilon}_{t+1})^2 \right) \right] = 0$. Hence, the new error term has again zero expectation under a variety of different possible assumptions, e.g. homoskedastic expectation errors across households and over time.¹⁷ As a consequence of the approximation, omitted higher-order moments are buried in the intercept. Although this makes it impossible to recover the individual time preference rate β , we can still consistently estimate all other preference parameters under the above assumptions.

Combining with our model for marginal utility, (3), we obtain

$$\begin{aligned} & \frac{1}{\sigma} \left(D_t^h \theta - \ln (C_t^h) + \gamma \left(ARITM \left[\ln (C_t^j) | X_t^j = X_t^h \right] \right) \right) \\ &= \left(\ln \beta - \frac{1}{2} s^2 \right) + \ln R_{t+1} \\ & \quad + \frac{1}{\sigma} \left(D_{t+1}^h \theta - \ln (C_{t+1}^h) + \gamma \left(ARITM \left[\ln (C_{t+1}^j) | X_{t+1}^j = X_{t+1}^h \right] \right) \right) + \zeta_{t+1}, \end{aligned} \quad (18)$$

¹⁷A more primitive condition implying the above restriction is to assume joint log-normality of the relevant variables. Although fairly restrictive, such an assumption is not uncommon in the literature on estimating linearized Euler equations. Alternatively and less restrictive, we would obtain a similar estimating equation by assuming that the innovations to the conditional moments of $\tilde{\varepsilon}_{t+1}$ are uncorrelated with the instruments used in the estimation.

or, rearranging terms,

$$\begin{aligned} \Delta \ln \left(C_{t+1}^h \right) &= \sigma \left(\ln \beta - \frac{1}{2} s^2 \right) + \Delta D_{t+1}^h \theta \\ &+ \gamma \left(ARITM \left[\ln \left(C_{t+1}^j \right) | X_{t+1}^j = X_{t+1}^h \right] - ARITM \left[\ln \left(C_t^j \right) | X_t^j = X_t^h \right] \right) \\ &+ \sigma \ln R_{t+1} + \sigma \zeta_{t+1}. \end{aligned} \quad (19)$$

In order to further simplify the above equation, we use the approximation $X_t^h \approx X_{t+1}^h$ to obtain

$$\begin{aligned} \Delta \ln \left(C_{t+1}^h \right) &= \sigma \left(\ln \beta - \frac{1}{2} s^2 \right) + \Delta D_{t+1}^h \theta + \gamma ARITM \left[\Delta \ln \left(C_{t+1}^j \right) | X_t^j = X_t^h \right] \\ &+ \sigma \ln R_{t+1} + \sigma \zeta_{t+1} + \varrho_{t+1}, \end{aligned} \quad (20)$$

where ϱ_{t+1} denotes the additional approximation error. While this last approximation is not required for estimation in a standard Euler equation framework, it will prove very convenient for the later derivation of further equilibrium conditions that will be needed to distinguish between true consumption externalities and merely correlated effects. Note also that the above approximation is very accurate, because all but one of the stratification variables we consider never change for a given household over time.¹⁸ In particular, for the predominant case in which strata in t and $t + 1$ are composed by the same households, the induced approximation error term is identically zero by construction. Moreover, even a more substantial approximation error would not pose a problem, as long as it is uncorrelated with the instruments used in the estimation.

Finally, introducing some short-hand notation for the intercept and error terms, we can represent the above linearized Euler equation by

$$\Delta \ln \left(C_{t+1}^h \right) = \alpha + \Delta D_{t+1}^h \theta + \gamma ARITM \left[\Delta \ln \left(C_{t+1}^j \right) | X_t^j = X_t^h \right] + \sigma \ln R_{t+1} + \varepsilon_{t+1}, \quad (21)$$

which coincides with equation (4) in the paper.

A3: Data Cleaning Procedures and Sample Selection

In total, our original sample contains 90,414 household year observations. We match household information across years by means of the history of interview numbers provided with each wave of the PSID. Where matching is not possible (for example, because information on interview numbers from past years is missing) or ambiguous, we drop the respective observations from the sample. In total, these deletions amount to a loss of 2,198 household years. We then proceed by cleaning our sample from observations with implausible, missing or topcoded information. Specifically, we delete 3,227 observations because of a zero in reported food expenditures at

¹⁸Age, gender, race and education (as defined for our purposes) are entirely time-invariant in our data. Thus, only our measure of "urbanity" displays some (very limited) temporal variation.

home, 92 because of topcoding of this variable, 27 because of topcoded food expenditures at restaurants, and 1,896 because of bad accuracy codes indicating that the respective information is poorly measured. Further, we only consider households whose head is between 18 and 65 years old. This restrictions leads to the deletion of another 10,094 household years.

Note that our data set is an unbalanced panel, because split-offs are treated as independent households from the moment of the split-off. Moreover, we also treat households with a head change as new households. The year in which the head change occurs is deleted from the sample (5,991 household years). We delete 344 observations for which there is no educational information, 288 observations because the household did not reside in continental USA, 60 household years because of missing race information and four because there is no major adult present in the respective household and year.

As the PSID was mainly designed to study the income dynamics of the poor, it significantly oversamples these households relative to the population. We follow the literature and leave the entire poverty subsample out of consideration. This means the loss of another 21,335 observations. Lastly, as we are interested in estimating a consumption Euler equation, we also have to discard households that are observed for a single year only (1,233 observations). In total, the above data cleaning procedures amount to a deletion of 44,972 household years, which, nevertheless, leaves us with a fairly large sample consisting of 43,244 observations. However, we must also discard observations with missing values for the individual-specific instruments used in the estimation. These deletions leave us with a sample of 26,358 observations. Lastly, in order to eliminate the influence of extreme outliers we also eliminate the observations with the 0.75% highest and lowest consumption growth rates leading to a loss of 396 observations.

After having constructed our peer groups as stated in the text, we still have to delete household years that are associated with groups consisting of less than 15 observations. Deletions related to small group sizes amount to a loss of 7,836, thus yielding a final sample of 18,126 observations, on which we conduct our analysis. Thus, although our necessary cleaning procedures substantially reduce the size of the data set, we are still left with a sufficiently big baseline sample of more than 18,000 observations to conduct meaningful inference.

A4: Econometric Issues

Estimation

As all of the above models are formulated within a forward-looking, rational expectations framework, they give rise to conditional moment restrictions that lend themselves to semiparametric estimation using GMM. Specifically, estimation is based on the orthogonality conditions implied by rational expectations, coupled with the standard assumption that instruments are uncorrelated with higher-order moments buried in the intercept, due to the linearization.¹⁹ Further, in this framework it is straightforward to account for the endogeneity of the households' after-tax interest rates as well as the relevant peer-group means in equations (5) and

¹⁹See, for example, Attanasio and Browning (1995), p. 1125.

(8), by excluding these variables from the instrument set and including other (lagged) variables instead. Hence, the starting point is a set of moment conditions of the form

$$E_t [u_{t+1}|z_t] = 0, \quad (22)$$

where u_{t+1} represents the error term from the respective Euler equation and z_t denotes the set of instrumental variables contained in the information set of period t . A necessary condition for identification is that the dimension of the instrument set be larger or equal to the number of parameters we want to estimate. Following the literature on Euler equation estimation, we transform the set of conditional moment restrictions into unconditional ones. Estimation thus exploits moment conditions of the following type:

$$E_t [u_{t+1}z_t] = 0. \quad (23)$$

Dropping time subscripts for notational convenience and specifying the determinants of the expectation error u , we can write

$$E [uz] = E [f(y, z, w, ARITM[\cdot|g(X)], \kappa)] = 0, \quad (24)$$

where $f(\cdot)$ now summarizes the dependence of the moment condition on the regressands y , the instruments z , the standard regressors w , the reference group means $ARITM[\cdot|g(X)]$ and the structural parameters κ , while $g(X_t^j)$ with $t = 1, \dots, T$ and $j = 1, \dots, G_t$ represents an identifier for the different reference groups by year and characteristics X_t^j . Lastly, $ARITM[\cdot|g(X)]$ is short-hand for the vector of all $ARITM[\cdot|g(X_t^j)]$.

Estimation of unconditional moment models of the form (24) is sufficiently standard. Because of the nonlinearities in (8), we use numerical optimization to obtain nonlinear GMM estimates. Throughout this paper, we only report results from two-stage estimation, noting that the results remain virtually unchanged if further iterations are carried out (iterated GMM). It is worthwhile to keep this in mind, as iterated GMM is normalization-invariant, while two-stage GMM is not. The fact that our results are not sensitive to the number of iterations also rules out the issue that Sargan tests may be adversely affected by differences in normalization.

The presence of generated regressors in both (5) and (8) leads to additional complications. Specifically, as reference group means $ARITM[\cdot|X_t^j = X_t^h]$ are estimated in a separate first step, we must account for the sampling variability associated with the respective estimates $ARITM_N[\cdot|g(X)]$ to conduct proper inference. This is of great importance, given that we would like to uncover possible peer effects at the population level rather than within our sample only.²⁰ The next section contains a brief discussion of how we make the required adjustments.

²⁰Most empirical studies of social interactions do not account for the presence of first-stage estimates when conducting inference. While such an approach may be sensible for the case of "local" interactions prevalent within a specific sample, e.g. neighborhood effects, it is clearly inadequate for studying large-group social effects based on a random sample of the whole population of interest.

Inference

To obtain consistent variance estimates in the presence of generated regressors, we follow Newey and McFadden (1994) and adopt a "joint GMM" interpretation for the two estimation steps. Basically, we "stack" the respective moment conditions from both estimation steps to form an extended vector of moments. The derivations are considerably simplified by formulating the moment conditions for both steps with reference to the entire sample rather than by individual reference group. Thus, let $d^h(X_t^j)$ with $t = 1, \dots, T$ and $j = 1, \dots, G_t$ denote a dummy equal to one at time t if household h has characteristics X_t^j and zero otherwise. For any given household h and time t , there is exactly one dummy equal to one, i.e. the dummy corresponding to her respective reference group $g(X_t^j)$. This step allows us to re-write the structural models (5) and (8) as

$$\begin{aligned} \Delta \ln(C_{t+1}^h) &= \alpha + \Delta D_{t+1}^h \theta + \sigma \ln R_{t+1} + X_t^h \lambda \\ &+ \gamma \sum_{t=1}^T \sum_{j=1}^{G_t} d^h(X_t^j) ARITM_N \left[\Delta \ln(C_{t+1}^j) | g(X_t^j) \right] + \varepsilon_{t+1} \end{aligned} \quad (25)$$

and, respectively,

$$\begin{aligned} \Delta \ln(C_{t+1}^h) &= \alpha + \Delta D_{t+1}^h \theta + \frac{\gamma}{1-\gamma} \sum_{t=1}^T \sum_{j=1}^{G_t} d^h(X_t^j) ARITM_N \left[\Delta D_{t+1}^h | g(X_t^j) \right] \theta \\ &+ \sigma \ln R_{t+1} + \frac{\gamma \sigma}{(1-\gamma)} \sum_{t=1}^T \sum_{j=1}^{G_t} d^h(X_t^j) ARITM_N \left[\ln R_{t+1} | g(X_t^j) \right] \\ &+ \left(1 + \frac{\gamma}{1-\gamma} \right) X_t^h \lambda + u_{t+1}. \end{aligned} \quad (26)$$

The number of generated regressors in each case is given by the number of estimated reference group means times the number of group years for which they have to be estimated. Thus, we have 553 and 2,212 generated regressors in the first and second case, respectively. The advantage of re-writing the model in this way is that we can now express the estimated reference group means $ARITM_N \left[\cdot | g(X_t^j) \right]$ as coefficients from first-step OLS regressions of the relevant variables v , i.e. $\Delta \ln(C_{t+1})$, ΔD_{t+1} and $\ln R_{t+1}$, on the group-year dummies $d(X_t^j)$ estimated on the entire sample. The additional moment conditions that will account for the sampling variability introduced by the first-step estimates are then nothing but the scores of these auxiliary OLS regressions. The structure of the scores is relatively simple: their components are the products of residuals and dummy variables, with all but one of the latter being identical zero by construction. Hence, many of the additional first-step moment conditions can be conveniently ignored. Moreover, since both first- and second-step moment conditions are defined for the full sample, corrected variance estimates can be computed using the techniques presented in Newey and McFadden (1994).

Specifically, for $t = 1, \dots, T$ and $j = 1, \dots, G_t$, let $d(X)$ denote the vector of all $d(X_t^j)$. Further, let $m(v, d(X), ARITM[\cdot|g(X)])$ denote the scores of the first-step regression generating the group averages of the relevant variables v . Then our first-step moment conditions are obviously given by

$$E[m(v, d(X), ARITM[\cdot|g(X)])] = 0. \quad (27)$$

Since the estimates for $ARITM[\cdot|g(X)]$ are used as (generated) regressors in the second step, we can now write the corresponding second-step moment conditions as

$$E[f(y, z, w, ARITM_N[\cdot|g(X)], \kappa)] = 0, \quad (28)$$

where the \sqrt{N} -consistent estimator for reference group means, $ARITM_N[\cdot|g(X)]$, has replaced the true population counterparts, $ARITM[\cdot|g(X)]$. Applying Newey and McFadden (1994, Theorem 6.1), we obtain an asymptotic distribution for the structural parameter estimates $\hat{\kappa}$ of the second stage given by

$$\sqrt{N * T}(\hat{\kappa} - \kappa) \sim N(0, V) \quad (29)$$

with

$$V = F_{\kappa}^{-1} E[\{f(\cdot) + F_{ARITM}\Psi(v, d(X))\} \{f(\cdot) + F_{ARITM}\Psi(v, d(X))\}'] F_{\kappa}^{-1'} \quad (30)$$

and

$$f(\cdot) = f(y, z, w, ARITM[\cdot|g(X)], \kappa) \quad (31)$$

$$F_{\kappa} = E[\nabla_{\kappa} f(y, z, w, ARITM[\cdot|g(X)], \kappa)] \quad (32)$$

$$F_{ARITM} = E[\nabla_{ARITM} f(y, z, w, ARITM[\cdot|g(X)], \kappa)] \quad (33)$$

$$M = E[\nabla_{ARITM} m(v, d(X), ARITM[\cdot|g(X)])] \quad (34)$$

$$\Psi(v, d(X)) = -M^{-1} m(v, d(X), ARITM[\cdot|g(X)]), \quad (35)$$

where ∇_{κ} and ∇_{ARITM} denote partial derivatives with respect to κ and $ARITM[\cdot|g(X)]$, respectively.

Each component of the adjusted variance matrix can be computed from its corresponding sample analog. Note that the adjustment for the presence of generated regressors is embodied in the expression $F_{ARITM}\Psi(v, d(X))$ in (30). It is fairly easy to check that for each variable over which we estimate peer-group means, this correction matrix amounts to the negative deviations of a household's own realizations from the respective peer-group means multiplied by the respective peer-group coefficients and the average realizations of the instruments among the peers. Thus, for the extreme case in which there are no social interactions at all, the correction terms become zero and the variance formula collapses to the standard one.²¹ Intuitively, sampling variability in the estimation of the reference group means is irrelevant for cases in

²¹In this case $F_{ARITM} = 0$ holds.

which there are no social interactions. On the other hand, the correction may become large depending on the estimated reference group coefficient as well as the sampling variability in the variables for which the means are estimated.

One final comment is in order: Since we use the estimated reference group means of the exogenous variables as both regressors and instruments, it might seem that further adjustments for the presence of generated instruments are required. This is not true, however, as all measurable functions of any predated variable provide valid instruments. Thus, given our conditional moment restrictions, the generated instruments have no effect on the asymptotic variance of the GMM estimator.²²

²²See Wooldridge (2002) p. 400 ff. for a more detailed discussion.

References

- ABEL, A. (1990). "Asset Prices Under Habit Formation and Catching Up With the Joneses", *American Economic Review*, 80, 38-42.
- ALESSIE, R. and A. KAPTEYN (1991). "Habit Formation, Interdependent Preferences and Demographic Effects in the Almost Ideal Demand System", *The Economic Journal*, 101, 404-419.
- ALESSIE, R. and F. TEPPA (2002). "Saving and Habit Formation: Evidence From Dutch Panel Data", Center Working Paper.
- ATTANASIO, O. (1999). "Consumption", in: J. Taylor and M. Woodford (eds.): *Handbook of Macroeconomics*, Volume I B, chapter 11. North Holland, Amsterdam.
- ATTANASIO, O., J. BANKS, C. MEGHIR, and G. WEBER (1999). "Humps and Bumps in Lifetime Consumption", *Journal of Business and Economics Statistics*, 17, 22-35.
- ATTANASIO, O. and M. BROWNING (1995). "Consumption over the Life Cycle and over the Business Cycle", *American Economic Review*, 85, 1118-1137.
- ATTANASIO, O. and G. WEBER (1993). "Consumption Growth, the Interest Rate and Aggregation", *Review of Economic Studies*, 60, 631-649.
- ATTANASIO, O. and G. WEBER (1995). "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey", *Journal of Political Economy*, 103, 1121-1157.
- BINDER, M. and H. PESARAN (2001). "Life-cycle Consumption Under Social Interactions", *Journal of Economic Dynamics and Control*, 25, 35-83.
- BLUNDELL, R., M. BROWNING, and C. MEGHIR (1994). "Consumer Demand and the Life-Cycle Allocation of Household Expenditures", *Review of Economic Studies*, 61, 57-80.
- BLUNDELL, R., L. PISTAFERRI, and I. PRESTON (2005). "Imputing Consumption in the PSID Using Food Demand Estimates from the CEX", IFS Working Paper 04/27.
- BROCK, W. and S. DURLAUF (2001). "Interactions-Based Models", in: J. Heckman and E. Leamer (eds.): *Handbook of Econometrics*, Vol. 5, 3297-3380.
- BROWNING, M. and M. COLLADO (2004). "Habits and Heterogeneity in Demand: A Panel Data Analysis", CAM Working Paper.
- BROWNING, M. and T. CROSSLEY (2001). "The Life Cycle Model of Consumption and Saving", *Journal of Economic Perspectives*, 15, 3-22.

- CAMPBELL, J. and J. COCHRANE (1999). "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior", *Journal of Political Economy*, 107, 205-251.
- COCHRANE, J. (1991): "A Simple Test of Consumption Insurance", *Journal of Political Economy*, 99, 957-976.
- DEATON, A. (1992). "Understanding Consumption", Oxford.
- DUESENBERY, J. (1949). "Income, Saving and the Theory of Consumer Behavior", Cambridge, MA.
- DYNAN, K. (2000). "Habit Formation in Consumer Preferences: Evidence from Panel Data", *American Economic Review*, 90, 391-406.
- FESTINGER, L. (1954). "A Theory of Social Comparison Processes", *Human Relations*, 117-140.
- GAERTNER, W. (1974). "A Dynamic Model of Interdependent Consumer Behavior", *Zeitschrift für Nationalökonomie*, 327-344.
- GALI, J. (1994). "Keeping Up with the Joneses: Consumption Externalities, Portfolio Choice and Asset Prices", *Journal of Money, Credit, and Banking*, 26, 1-8.
- GUARIGLIA, A. and M. ROSSI (2002). "Consumption, Habit Formation and Precautionary Saving: Evidence from the British Household Panel Survey", *Oxford Economic Papers*, 54, 1-19.
- HALL, R. and F. MISHKIN (1982). "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households", *Econometrica*, 50, 461-481.
- JACOBS, K. (1999). "Incomplete Markets and Security Prices: Do Asset-Pricing Puzzles Result from Aggregation Problems?", *Journal of Finance*, 54, 123-164.
- KAPTEYN, A., S. VAN DER GEER, H. VAN DE STADT, and T. WANSBEEK (1997). "Interdependent Preferences: An Econometric Analysis", *Journal of Applied Econometrics*, 12, 665-686.
- LAWRANCE, E. (1991). "Poverty and the Rate of Time Preference: Evidence from Panel Data", *Journal of Political Economy*, 99, 54-77.
- LEIBENSTEIN, H. (1950). "Bandwagon, Snob, and Veblen Effects in the Theory of Consumers' Demand", *Quarterly Journal of Economics*, 64, 183-207.
- LJUNGQVIST, L. and H. UHLIG (2000). "Tax Policy and Aggregate Demand Management under Catching Up with the Joneses", *American Economic Review*, 90, 356-366.

- LUTTMER, E. (2004). "Neighbors as Negatives: Relative Earnings and Well-being", NBER Working Paper 10667.
- MACE, B. (1991): "Full Insurance in the Presence of Aggregate Uncertainty", *Journal of Political Economy*, 99, 928-956.
- MANSKI, C. (1993). "Identification of Endogenous Social Effects: The Reflection Problem", *Review of Economic Studies*, 60, 531-42.
- MANSKI, C. (1995). "Identification Problems in the Social Sciences", Cambridge.
- MANSKI, C. (2000). "Economic Analysis of Social Interactions", *Journal of Economic Perspectives*, 14, 115-136.
- NEWKEY, W. and D. MCFADDEN (1994). "Large Sample Estimation and Hypothesis Testing", in: Engle, R. and D. McFadden (eds.): *Handbook of Econometrics*, Vol. 4, Ch. 36, 2111-2245.
- POLLAK, R. (1976). "Interdependent Preferences", *American Economic Review*, 309-320.
- RUNKLE, D. (1991). "Liquidity Constraints and the Permanent-Income Hypothesis", *Journal of Monetary Economics*, 27, 73-98.
- SKINNER, J. (1987). "A Superior Measure of Consumption from the Panel Study of Income Dynamics", *Economics Letters*, 23, 213-216.
- SOULELES, N. (2005). "Consumer Sentiment: Its Rationality and Usefulness in Forecasting Expenditure – Evidence from the Michigan Micro Data", *Journal of Money, Credit and Banking*, forthcoming.
- VAN PRAAG, B. and P. FRIJTERS (1999). "The Measurement of Welfare and Well-being: The Leyden Approach", in: D. Kahneman, E. Diener and N. Schwarz (eds.): *Well-being: the Foundations of Hedonic Psychology*, New York, 413-432.
- WOITTIEZ, I. and A. KAPTEYN (1998). "Social Interactions and Habit Formation in a Model of Female Labour Supply", *Journal of Public Economics*, 70, 185-205.
- WOOLDRIDGE, J. (2002). "Econometric Analysis of Cross Section and Panel Data", Cambridge, MA.
- ZELDES, S. (1989). "Consumption and Liquidity Constraints: An Empirical Analysis", *Journal of Political Economy*, 97, 305-346.