

**SMOOTH IT LIKE THE “JONESES?”
ESTIMATING PEER-GROUP EFFECTS IN
INTERTEMPORAL CONSUMPTION CHOICE**

Jürgen Maurer and André Meier

167-2008

Smooth it Like the "Joneses?" Estimating Peer-Group Effects in Intertemporal Consumption Choice*

Jürgen Maurer[†] and André Meier[‡]

This Version: September 2007

Abstract

Recent theoretical contributions have suggested peer-group effects as a potential explanation for several puzzles in macroeconomics, but their empirical relevance for intertemporal consumption choice is an open question. We derive an extension of the standard life-cycle model that allows for consumption externalities. In this framework, we propose a social multiplier approach to distinguish true externalities from merely correlated effects. Estimating our model using US panel data, we find strong predictable co-movement of household consumption within peer groups. Although much of this co-movement reflects correlated effects only, there is statistically significant evidence for moderate consumption externalities across several plausible peer-group specifications.

JEL classification: C23, D12, D91, Z13

*A shorter version of this paper appeared in *The Economic Journal*, Vol. 118 (527), pp. 454–476

[†]Mannheim Research Institute for the Economics of Aging (MEA), University of Mannheim, Germany. Maurer gratefully acknowledges the hospitality of the Institute for Fiscal Studies in London as well as financial support provided by the European Community's Human Potential Programme through contract HPRN-CT-2002-00235. Address for correspondence: Jürgen Maurer, Mannheim Research Institute for the Economics of Aging (MEA), University of Mannheim, L13, 17, 68131 Mannheim, Germany. E-mail: maurer@mea.uni-mannheim.de.

[‡]International Monetary Fund, 700 19th Street, Washington DC, 20431. E-mail: ameier@imf.org. The views expressed in this paper are those of the authors and should not be interpreted as the views of the International Monetary Fund. 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 two anonymous referees, Orazio Attanasio, Richard Blundell, Steven Bond, Martin Browning, Tullio Jappelli, Søren Leth-Petersen, 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. A previous version of this paper, Maurer and Meier (2005), was circulated under the title: "Do the "Joneses" Really Matter? Peer-Group vs. Correlated Effects in Intertemporal Consumption Choice". All remaining errors are our own responsibility.

Consumption is arguably a social experience, and the position of other people with respect to our own consumption often matters to us. The idea of social interaction is reflected, for example, in references to "conspicuous consumption" or "keeping up with the Joneses", which are commonplace in casual discussions about the determinants of particular consumption patterns. Consistent with this, psychologists, sociologists and economists have gathered evidence showing that consumers' well-being is affected by their *relative* economic standing rather than their absolute resources alone.¹ Within economics, early discussions of consumption externalities, or peer effects, actually date back at least to the seminal contributions of Duesenberry (1949) and Leibenstein (1950). More recently, a growing literature in macroeconomics and finance has resorted to models featuring consumption externalities in order to explain prominent empirical puzzles.²

Against this backdrop, it is surprising that economists have largely ignored consumption externalities when studying the determinants of intertemporal consumption choice in actual micro-level data. Clearly, some "stylized facts" about household consumption are suggestive of, or at least compatible with, peer effects: rather than being smooth, life-cycle consumption profiles feature humps and bumps whose exact shapes appear to depend on characteristics of the respective households. On the one hand, this might simply mirror changing demographic situations that have a direct effect on the utility derived from a given level of consumption. On the other hand, synchronized consumption patterns within groups could also indicate the presence of peer effects. Given the empirical support for peer-group phenomena in other fields, their relevance for intertemporal consumption decisions surely deserves closer investigation.

We begin our own study by proposing a theoretical model of consumption that allows for peer effects. Instead of introducing an ad hoc behavioural 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 fully consistent with the usual forward-looking utility maximization framework as presented, for example, in Browning and Crossley (2001). Apart from providing a coherent economic foundation for our exercise, this approach also ensures

¹A recent example is Luttmer (2006), who also provides further references to this quite sizable literature.

²Prominent examples of this line of research include Abel (1990), Galí (1994), Campbell and Cochrane (1999), Ljungqvist and Uhlig (2000) or Binder and Pesaran (2001).

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.

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 estimation techniques uncover substantial predictable consumption co-movement within peer groups, suggesting the presence of strong consumption externalities. However, these results have to be taken with caution, because estimation under the usual Euler equation framework is vulnerable to even minor misspecifications. The difficulty 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 really affected by peer-group behaviour (endogenous effects) or does it display co-movement within groups merely because individuals share similar unobserved characteristics or suffer similar predictable shocks (correlated effects)?

Disentangling actual peer-group phenomena from correlated effects constitutes the principal challenge for our empirical application. 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; using these, we can re-assess peer effects in intertemporal consumption and provide a more robust test of their relevance. Interestingly, once correlated effects are adequately accounted for, our estimation results no longer point to the very strong consumption externalities suggested by the standard estimation approach. Still, we find statistically significant evidence for moderate peer effects across a range of plausible specifications.

Apart from the aforementioned early contributions that have inspired this paper, our research is also related to a number of papers in the more recent literature. First, our analysis extends existing microeconomic studies of intertemporal consumption profiles by incorporating peer effects. Previous research has shown that a simple power utility version

of the intertemporal model cannot explain key features of life-cycle consumption. However, some progress has been achieved with the inclusion of demographic preference shifters, as suggested, for example, by Blundell et al. (1994) and Attanasio et al. (1999). Because our approach nests this class of models, we can assess the importance of peer 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, judging from the papers by Dynan (2000), Guariglia and Rossi (2002), Alessie and Teppa (2002), Ravina (2005) and Browning and Collado (2007). Our approach differs from these contributions in that we focus on "external" habit formation: instead of 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. Our work is, therefore, a natural complement to the literature on internal habits, notably Dynan (2000), who estimates modified Euler equations based on the same US micro data that we use. Most closely related to our work is the study by Ravina (2005). Using expenditure data from Californian credit card holders, she studies the role of internal and external habits for intertemporal consumption choice. However, her paper defines "external habits" based on a strictly geographic notion of "neighbourhood effects", whereas we focus on peer groups defined in terms of sociodemographic characteristics.

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. Essentially, these authors investigate the role of "fads and fashions" in the allocation of total expenditure to certain categories of consumption, such as clothing, food or transportation. In terms of economic analysis, 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, namely the intertemporal allocation decision.

Third, our analysis sheds light on the empirical plausibility of recent theoretical con-

tributions that have suggested consumption externalities as potential solutions to empirical puzzles in macroeconomics and finance. Indeed, the recent macroeconomic literature has readily adopted various kinds of internal and external habit specifications, although supporting evidence from microeconomic studies is either scarce or absent.

The remainder of this paper is organized as follows. Section 1 is preparatory and reviews the standard life-cycle model, which we extend, in section 2, to allow for peer effects. In section 3, we present our data and describe the way we construct peer groups. This is followed by a detailed exposition of econometric issues in section 4, where we derive our model specification and discuss identification and inference. In section 5, we turn to our results and their interpretation. Section 6 concludes with some final remarks. Less instructive derivations and econometric technicalities are relegated to the appendix.

1 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 a finite-lived consumer h at time t who faces a riskless asset with real after-tax rate of return R_{t+1}^h . Assume that the consumer has von Neumann-Morgenstern preferences and derives utility from consumption C^h , 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

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

subject to an intertemporal budget constraint; E_t denotes the conditional expectations operator in time t , $T - t$ the consumer's remaining life-span. The first-order condition for this problem is the familiar Euler equation

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

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, with the after-tax interest rate R_{t+1}^h representing the opportunity cost of period t consumption relative to consumption in period $t + 1$.

2 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 demographic preference shifters. To illustrate this point, we show in appendix A1 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. Thus, apart from the key determinants of marginal utility already considered in their model, we 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, such as family size or the number of children. Such preference shifters can also be motivated as implicit equivalence scales. Accordingly, the specification in (3) flexibly allows for time-variant preferences as well as implicit adjustments of household consumption in response to changing family compositions.

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 that share some common sociodemographic 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, capturing notions of status concern or jealousy. Thus, our specification encompasses intertemporal consumption complementarities between similar households. The central parameter of interest is γ , which captures the strength of possible peer effects. It measures the extent to which the consumption of a given household reacts to the average household consumption of its respective peers.³

Although it might first seem somewhat ad hoc, the specification given by (3) is really but a slight extension of the standard power utility model. In particular, 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 obtain

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

³Note 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 differences in household size. However, since it is a priori unclear which of the two approaches is more plausible, we retain the assumption that any relevant consumption externalities operate at the household level. This formulation has the advantage that changes in household composition provide an extra source of within-peer-group variation in intertemporal consumption, thus facilitating inference.

where α contains both the logarithm of the discount rate, β , and higher-order terms stemming from the linearization.⁴ 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 test 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. For example, one might argue that different degrees of education—an important dimension for social comparisons—might also imply different degrees of impatience, i.e., higher or lower discount rates β . As such a phenomenon would concern the whole peer group, similar behaviour 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

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

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 literature 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.⁵ Thus, by including the additional term $X_t^h \lambda$, we implicitly parameterize $\ln \beta$, which was previously buried in the intercept.⁶

Our new, augmented Euler equations reads as

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

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. 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, $ARITM_N \left[\Delta \ln \left(C_{t+1}^j \right) | X_{t+1}^j = X_t^h \right]$, gives rise to a generated regressor problem that we need to take into account when computing asymptotic standard errors.

Finally, 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 of providing estimates of the relevant preference parameters without a full characterization of the particular equilibrium. Even so, identification in this context requires that there exist sufficient exclusion restrictions to separate individual-level determinants of consumption growth from potential peer effects. In our case, this requirement is met by within-peer-group variation in demographic preference shifters and/or the after-tax interest rate. Specifically, the demographic preference shifters exploit the fact that each household aims at smoothing expected discounted marginal utility, which is

⁵Prominent examples include Lawrance (1991), Attanasio and Browning (1995) and Dynan (2000).

⁶A different issue is the likely presence of common unpredictable shocks within peer groups. It is important to note that such shocks do not confound our analysis, because peer-group consumption—an endogenous variable—will be instrumented with lagged information throughout. In this sense, we are only dealing with predictable or planned co-movement of consumption that is picked up by our instruments.

affected by its own demographic composition. This gives rise to exogenous within-peer-group variation in household consumption levels. Heterogeneity in after-tax interest rates, in turn, implies differences in the intertemporal substitution motive within peer groups that provide an additional source of independent variation. Conceptually, these two features also distinguish the model in (5) from tests of full consumption insurance with perfect capital markets and pareto-efficient equilibria which are based on marginal utility rather than consumption levels (see, for example, Mace (1991) and Cochrane (1991)).

3 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 samples used in the estimation.

3.1. 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. While the PSID has a very broad content, including economic and demographic as well as sociological and psychological measures, its coverage of consumption behaviour 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 between 1974 and 1987. Accordingly, the PSID offers a maximum of 14 consecutive annual observations to investigate households' intertemporal consumption patterns.

Having data on food expenditure only is an obvious drawback, even if the empirical consumption literature has commonly used such data to explore consumer behaviour.⁷ The critical question is whether or not food consumption provides a reasonable proxy for overall

⁷Prominent examples include Hall and Mishkin (1982), Hotz et al. (1988), Zeldes (1989), Altug and Miller (1990), Lawrance (1991), Cochrane (1991), Runkle (1991), Shea (1995), Hayashi et al. (1996), Jacobs (1999), Dynan (2000), Cox et al. (2004), Hurst (2004) or Hurst and Stafford (2004).

consumption for the purposes of our particular application. A priori, we would expect other consumption categories, such as clothing or cars, to be more promising candidates for peer effects. However, it is important to keep in mind that we are not looking for intertemporal peer effects in specific consumption categories; instead we would like to assess their relevance for comprehensive consumption baskets. Accordingly, we wonder to what extent food consumption is representative for broader consumption measures in terms of conspicuousness. While there is not much evidence on this issue, recent work by Heffetz (2004) provides some reassurance. Based on his own survey, Heffetz constructs three different "visibility indices" to measure the conspicuousness of thirty-one distinct consumption categories covering more than 95% of consumption reported by the Consumer Expenditure Survey (CEX). In all three indices, both food at home and food outside home consistently rank in the (upper) middle of the distribution, with their exact positions ranging between seven and fifteen. The two food consumption categories continue to appear fairly representative once the expenditure categories are appropriately value weighted. Specifically, although some large categories, like clothing, cars or furniture, always rank among the top five categories on the list, other quantitatively important categories, such as utilities, medical care or insurance policies, are found much further down the visibility scales. Our tentative conclusion is that food consumption, while narrower an aggregate than we would wish, may not be a bad proxy for overall consumption in terms of susceptibility to peer effects.

As mentioned, the data we use cover the interview period 1974 through 1987. After necessary data cleaning, which is duly documented in appendix A3, our final data set comprises some 26,000 observations.

3.2. Peer-Group Construction

The specification of peer groups is critical for any analysis of social interactions. Because the reference groups we choose to consider are taken to be characteristic of a household's social environment, we must take a stance on what personal attributes plausibly define such an environment. The ideal solution would be to use observed behaviour 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)). In light of this, 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) 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 households whose heads 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, such as race, gender, family status, occupation and urbanity. Hence, we construct reference groups based on different subsets of the following attributes: age cohort, race, gender, the presence of children, educational attainment, occupational status and "size of the nearest city" as a measure for urbanity.⁸ 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 dummy indicating whether or not there are children living in the household; 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; a variable summarizing occupational status as either "blue collar", "white collar" or "other", based on the PSID's detailed employment information; 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 be extended even further. In a sample of limited size, however, this has to be traded off against the disadvantage of ending

⁸Some of these same variables have also been found to be important predictors for individual welfare functions over income as studied in the Leyden approach (see, for example, van Praag and Frijters, 1999), which provides additional support for our stratification strategy. We also experimented with regional dummies as both an alternative and a complement to "urbanity", but found the results essentially unaffected.

up with overly small reference groups or substantial data losses. Indeed, considering all of the seven above-mentioned stratification variables simultaneously already leads to a considerable reduction in sample size. We, therefore, examine a range of different stratifications, applying between five and seven of the above-mentioned characteristics at a time.

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. Even so, the need to delete observations pertaining to overly small reference groups causes a reduction of our original sample to final sizes between 18,126 and 8,324 observations, depending on the number and type of stratification variables under consideration.

Table 1 presents basic summary statistics for our samples before and after the deletion of households belonging to small peer groups. Comparison shows that our samples generally become "more male", "more white", "less urban", "more educated" and "more white collar" as well as somewhat "less childless" as a consequence of data deletions, whereas compositional changes in terms of birth cohorts are moderate. Essentially, imposing a lower bound on cell size removes most of the observations pertaining to less common household heads. Although this is of course unfortunate, we still find most of our final samples relatively well-balanced even with respect to other studies that do not face data constraints associated with peer-group construction. For example, many authors in the consumption literature have dropped female-headed households from their sample at the outset, 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.

Table 1: Sample summary statistics before and after deletion of small cells*

	Before	Samples after deletion of small cells					
	deletions	(1)	(2)	(3)	(4)	(5)	(6)
Cohort (age in 1974)							
5-10	0.03	0.00	0.00	0.00	0.00	0.00	0.00
11-16	2.54	1.71	1.95	1.09	1.28	1.38	0.80
17-22	15.91	17.00	17.58	18.56	17.61	19.57	21.06
23-28	23.20	26.24	26.59	31.17	29.71	32.71	41.76
29-34	13.59	14.98	14.97	16.94	15.49	16.49	19.31
35-40	10.11	9.10	9.84	7.52	8.87	7.79	6.86
41-46	13.02	13.50	12.54	10.12	11.81	12.20	6.74
47-52	11.50	12.19	11.26	10.03	10.49	8.27	3.29
53-58	7.35	4.18	4.39	3.23	3.87	1.22	0.00
59-65	2.75	1.10	0.86	1.32	0.86	0.37	0.18
Gender							
female	13.89	0.42	0.64	0.51	0.50	0.55	0.41
male	86.11	99.58	99.36	99.48	99.50	99.45	99.59
Race							
white	90.56	99.91	100.00	100.00	100.00	100.00	100.00
non-white	9.44	0.09	0.00	0.00	0.00	0.00	0.00
City Size							
less than 50,000	39.55	45.44	44.21	44.78	46.32	46.68	46.40
between 50,000 and 500,000	38.08	38.09	39.04	39.59	39.20	37.81	38.35
more than 500,000	22.37	16.46	16.76	15.63	14.48	15.51	15.26
Education							
less than high school	21.79	14.67	17.69	12.14	17.24	11.12	6.33
high school or more	36.33	36.89	35.51	36.06	35.26	31.18	29.32
finished college or more	41.88	48.44	46.80	51.81	47.50	57.69	64.34
Occupation							
"blue collar"	41.67	42.04	45.01	40.78	43.66	43.02	39.01
"white collar"	46.29	50.51	51.92	52.09	53.10	56.47	60.59
other	12.04	7.44	3.06	7.14	3.24	0.51	0.38
Any children							
no	42.80	37.40	35.93	30.83	30.11	31.57	16.59
yes	57.20	62.60	64.07	69.17	69.89	68.43	83.41
Sample Size							
	26,358	18,126	18,206	13,704	14,791	12,450	8,324
Stratification							
Age	...	X	X	X	X	X	X
Gender	...	X	X	X	X	X	X
Race	...	X	X	X	X	X	X
City size	...	X	X	X	X	X	X
Education	...	X		X		X	X
Occupation	...		X		X	X	X
Any children	...			X	X		X

*All figures are in percent. Computations for baseline samples (minimum cell size = 15).

4 Specification, Identification and Inference

4.1. 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. 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 has 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 thus 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 . 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.

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

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

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.

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

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

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

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

where $u_{t+1}^h \equiv \varepsilon_{t+1}^h + \frac{1}{1-\gamma} ARITM \left[\varepsilon_{t+1}^j | X_t^j = X_t^h \right]$ has been introduced as a convenient shorthand notation for the new combined error term. Equation (8) will serve as our principal estimating equation. Of course, in the estimation we need to replace the peer-group means $ARITM \left[\cdot | X_t^j = X_t^h \right]$ with the nonparametric estimates $ARITM_N \left[\cdot | X_t^j = X_t^h \right]$ based on the respective 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 it to spuriously pick 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, consider the interest rate—a theoretically undisputed determinant of consumption growth. To the extent that higher average interest rates faced by her peers raise average consumption growth in an individual’s peer group, they also cause the individual herself to raise consumption growth if peer-group effects are present. This effect operates in addition to any direct intertemporal substitution motive (as determined by the individual’s own interest rate) and explains why the coefficient in front of the peer-group interest rate contains $\frac{\gamma}{1-\gamma}$ as a multiplier. Indeed, equation (8) in principle allows disentangling the direct individual from the indirect social effects of all standard explanatory variables. The requirement is that there is some household-level variation in these variables relative to their respective peer-group averages. Consider again the example of interest rates. Many authors have estimated consumption Euler equations using pre-tax interest rates. We do not follow this practice here, precisely because it would preclude a distinction 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}^h as regressors, thereby identifying both effects of interest. Note further that 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.⁹ These variables (and their respective peer-group averages) need to be instrumented, which reduces the necessary variation even further. Hence, it may be

⁹By contrast, we maintain the standard assumption that the demographic preference shifters ΔD_{t+1} are exogenous to the consumption decision (both at the level of the individual and the peer group) and thus need not be instrumented.

difficult to obtain very precise estimates of the parameters of interest.

4.2. Explanatory Variables and Instrument Choice

4.2.1. *Explanatory variables.* In order to estimate Euler equations on micro data, it is essential to properly account for real-life heterogeneity and demographic variation. At a minimum, most economists have considered changes in household size as an important determinant of consumption growth.¹⁰ 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 variables in equation (5). Hence, X_t^h comprises all of the cohort, race, gender, children, education, occupation and urbanity dummies used for peer-group stratification in the respective sample. As a proxy for the riskless asset return R_{t+1}^h , we consider 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).

4.2.2. *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

¹⁰Prominent 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.

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 labour market status of head and spouse as instruments.

5 Estimation Results

5.1. Results for Equation (5)

We start by estimating (5) using two-step GMM. The exact formulation of the estimator is described in appendix A4, which 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 2 presents point estimates along with appropriately adjusted standard errors for all six peer-group stratifications. 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 results also differ little across peer-group specifications, although different stratification strategies obviously imply significant variation in sample size and composition. 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 between 0.06 and 0.07.

Table 2: Estimates for Euler equation including peer-group consumption growth

	(1)	(2)	(3)	(4)	(5)	(6)
Parameter						
Δ family size	0.091 (0.0063)	0.094 (0.0064)	0.099 (0.0089)	0.090 (0.0075)	0.105 (0.0082)	0.106 (0.0120)
Δ major adults	0.103 (0.0172)	0.109 (0.0176)	0.089 (0.0217)	0.131 (0.0210)	0.106 (0.0207)	0.107 (0.0291)
Δ children	-0.018 (0.0073)	-0.021 (0.0074)	-0.032 (0.0096)	-0.029 (0.0086)	-0.033 (0.0096)	-0.052 (0.0133)
σ	0.061 (0.0139)	0.058 (0.0141)	0.068 (0.0153)	0.055 (0.0145)	0.067 (0.0190)	0.069 (0.0207)
γ	0.955 (0.0270)	0.957 (0.0242)	0.943 (0.0298)	0.946 (0.0219)	0.883 (0.0443)	0.900 (0.0405)
Stratification						
Age	X	X	X	X	X	X
Gender	X	X	X	X	X	X
Race	X	X	X	X	X	X
City size	X	X	X	X	X	X
Education	X		X		X	X
Occupation		X		X	X	X
Any children			X	X		X
J-statistic	9.23	14.46	9.19	12.09	8.71	7.30
p-value	0.7556	0.3422	0.7584	0.5203	0.7943	0.8861
Sample size	18,126	18,206	13,704	14,791	12,450	8,324

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.

Most importantly, however, the parameter associated with peer-group consumption growth, γ , indicates strong consumption externalities across all stratifications, with point estimates between 0.88 and 0.96 and small standard errors. In addition, the usual specification tests lend support to these results in that the J-statistics clearly fail to reject the model at any conventional level of significance.

Basically, the results in Table 2 indicate that the instruments we use are orthogonal

to deviations of individual consumption growth from its reference-group means, controlling for other typical regressors. In other words, while consumption within (differently defined) 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 behaviour, 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.

5.2. Results for Equation (8)

Therefore we next turn to the estimation of (8). In each case, the regression again includes the full set of demographic control variables used for the respective stratification of peer groups. In addition, all of the estimates in Table 3 account for the presence of generated regressors.

Note first that the estimated effects of the demographic preference shifters 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 point estimates between 0.09 and 0.15, more in line with existing results from the literature. Peer-group effects, in turn, appear to be considerably *smaller* now: despite some variation across stratifications, the parameter estimates all lie below 0.45 and have sufficiently small standard errors to reject the very high values for γ reported in Table 2 at conventional levels of statistical significance. This, together with the much higher J-statistics in Table 3, clearly supports the hypothesis that inference about peer effects is vulnerable to the omission of (predictable) group-specific effects.

Table 3: Euler equation estimates imposing the social equilibrium condition

	(1)	(2)	(3)	(4)	(5)	(6)
Parameter						
Δ family size	0.096 (0.0065)	0.095 (0.0064)	0.101 (0.0086)	0.092 (0.0076)	0.109 (0.0082)	0.107 (0.0122)
Δ major adults	0.103 (0.0183)	0.109 (0.0184)	0.094 (0.0227)	0.140 (0.0219)	0.101 (0.0212)	0.117 (0.0302)
Δ children	-0.019 (0.0075)	-0.021 (0.0074)	-0.032 (0.0096)	-0.030 (0.0086)	-0.033 (0.0096)	-0.047 (0.0135)
σ	0.150 (0.0691)	0.134 (0.0592)	0.115 (0.0635)	0.130 (0.0539)	0.085 (0.0581)	0.093 (0.0628)
γ	0.111 (0.1803)	0.237 (0.1373)	0.296 (0.1255)	0.375 (0.0954)	0.367 (0.0929)	0.444 (0.1020)
Stratification						
Age	X	X	X	X	X	X
Gender	X	X	X	X	X	X
Race	X	X	X	X	X	X
City size	X	X	X	X	X	X
Education	X		X		X	X
Occupation		X		X	X	X
Any children			X	X		X
J-statistic	46.82	62.20	32.93	58.73	22.15	18.07
p-value	0.0001	0.0000	0.0076	0.0000	0.1383	0.3201
Sample size	18,126	18,206	13,704	14,791	12,450	8,324

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.

In the case of equation (5), such omitted effects systematically bias estimates of γ toward 1 while reducing the power of the J-test.

Although we can reject the very strong peer effects implied by our initial estimates based on (5), the results for certain stratifications still point to moderate consumption externalities. Columns (3) through (6) of Table 3, in particular, report estimates of γ that are significantly different from zero. Falling into a range from 0.30 to 0.44, these estimates are even fairly close

in size to the "neighbourhood effects" documented by Ravina (2005).¹¹ Thus, some evidence for economically meaningful peer effects persists even as we apply a more robust estimation approach based on social multipliers.

Given the range of estimates of γ shown in Table 3, it bears repeating that the estimated coefficients cannot be used to judge the particular peer-group specification under consideration: to attain identification, we have to either assume a certain intensity of peer effects and then try to learn about the particular groups in which these effects operate; or specify peer groups a priori and estimate the intensity of consumption externalities conditional on a given specification. This paper applies the latter approach. Accordingly, we cannot infer from Table 3 which of the different stratifications is most appropriate. However, readers with clear a priori views on relevant peer group characteristics will certainly attach greater weight to those columns in Table 3 that pertain to the corresponding stratification. At the same time, stratification choices cannot, of course, account for the significant difference in results between Tables 2 and 3, given that the estimated strength of peer effects is reduced by more than half across all stratifications.

Before addressing the implications of our findings in greater detail, it seems appropriate to supplement the statistical J-test of our model with a more "economic" test. For this purpose, Table 4 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 across all specifications, 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.

¹¹Ravina (2005) estimates an effect of 0.29 of city-level per capita taxable sales on individual credit card expenditure. Interestingly, this similarity of results should in principle be expected if geographical segregation based on sociodemographic characteristics is relatively widespread.

Table 4: Excess sensitivity tests

	(1)	(2)	(3)	(4)	(5)	(6)
Parameter						
Δ family size	0.095 (0.0066)	0.095 (0.0064)	0.100 (0.0087)	0.092 (0.0076)	0.108 (0.0084)	0.109 (0.0126)
Δ major adults	0.106 (0.0192)	0.111 (0.0192)	0.097 (0.0243)	0.140 (0.0232)	0.107 (0.0233)	0.128 (0.0319)
Δ children	-0.019 (0.0075)	-0.021 (0.0074)	-0.031 (0.0097)	-0.029 (0.0086)	-0.033 (0.0097)	-0.049 (0.0142)
σ	0.153 (0.0702)	0.136 (0.0593)	0.116 (0.0641)	0.133 (0.0544)	0.092 (0.0589)	0.089 (0.0623)
γ	0.101 (0.1841)	0.230 (0.1397)	0.291 (0.1271)	0.371 (0.0968)	0.365 (0.0936)	0.457 (0.0985)
Income growth	-0.012 (0.0264)	-0.011 (0.0285)	-0.012 (0.0314)	0.010 (0.0398)	-0.026 (0.0392)	-0.052 (0.0477)
Stratification						
Age	X	X	X	X	X	X
Gender	X	X	X	X	X	X
Race	X	X	X	X	X	X
City size	X	X	X	X	X	X
Education	X		X		X	X
Occupation		X		X	X	X
Any children			X	X		X
J-statistic	46.18	61.57	32.19	59.04	21.02	15.56
p-value	0.0000	0.0000	0.0061	0.0000	0.1361	0.4116
Sample size	18,126	18,206	13,704	14,791	12,450	8,324

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.

Furthermore, Table 5 checks the robustness of our results with respect to different minimum group sizes. Specifically, we re-estimate all specifications shown in Table 3 based on data with minimum cell sizes of 10, 20 and 25, respectively. The point estimates for the demographic preference shifters and the interest rate coefficient are essentially unaffected. Thus, to conserve space, the table only reports the results for γ . These also fall mostly into the same range as

before, although there appears to be a tendency for somewhat higher point estimates in smaller, more selective and homogeneous samples, be it because of a finer stratification strategy (as noted before) or because of a higher minimum cell size.

Table 5: Robustness checks with respect to changes in minimum cell sizes

	(1)	(2)	(3)	(4)	(5)	(6)
Min. Cell Size 10						
Y	0.061 (0.1861)	0.202 (0.1358)	0.126 (0.1255)	0.186 (0.1279)	0.285 (0.0928)	0.416 (0.0813)
J-statistic	55.82	65.87	52.10	61.03	42.83	34.40
p-value	0.0000	0.0000	0.0000	0.0000	0.0003	0.0048
Sample size	19,843	19,795	17,127	17,601	15,252	11,654
Min. Cell Size 20						
Y	0.163 (0.1801)	0.233 (0.1505)	0.361 (0.1458)	0.379 (0.1148)	0.436 (0.1016)	0.478 (0.1331)
J-statistic	38.71	59.30	22.89	43.16	14.69	21.09
p-value	0.0012	0.0000	0.1166	0.0003	0.5475	0.1749
Sample size	16,179	17,028	10,763	12,517	9,386	5,932
Min. Cell Size 25						
Y	-0.054 (0.3330)	0.139 (0.2004)	0.390 (0.1644)	0.424 (0.1493)	0.402 (0.1445)	0.594 (0.1282)
J-statistic	28.27	47.07	20.63	34.10	16.77	16.87
p-value	0.0294	0.0001	0.1932	0.0053	0.4005	0.3942
Sample size	13,333	15,717	8,191	9,742	7,217	4,247
Stratification						
Age	X	X	X	X	X	X
Gender	X	X	X	X	X	X
Race	X	X	X	X	X	X
City size	X	X	X	X	X	X
Education	X		X		X	X
Occupation		X		X	X	X
Any children			X	X		X

Estimates account for the presence of generated regressors. Standard errors are in parentheses.

The estimation also includes all other controls of the corresponding models presented in Table 3.

5.3. Discussion

In sum, our estimation results indicate moderate peer effects in intertemporal consumption for several plausible peer-group specifications but do not confirm the *prima facie* evidence in favour of very strong consumption externalities. Certainly, estimates from the simple Euler equation (5) reveal that there is strong predictable co-movement of consumption within peer groups. Once the analysis is cast in the more robust framework of Manski's reflection problem, however, this co-movement turns out to be largely driven by omitted group-specific factors rather than true consumption externalities: although the results for equation (8) point to the presence of moderate peer effects for certain stratifications, the estimated values for γ are consistently below 0.5 and significantly different from the high point estimates obtained from the simple Euler equation.

In this context, the large J-statistics for many specifications in Tables 3 through 5 are actually instructive in that they help to explain the apparent contradiction with the results from our initial regressions in Table 2. The contradictory results are, in fact, two sides of the same coin: If the estimated peer-group effects in Table 2 are chiefly 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.¹²

Even the moderate consumption externalities that we find for certain peer group specifications strike us as a remarkable finding. After all, the intuitively appealing notion of peer effects has a very specific and perhaps quite restrictive interpretation in our intertemporal consumption model: faced with consumption externalities, individuals will try to smooth their own consumption profile relative to that of their peers. This prediction may be quite distinct

¹²It 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 several of our results imply a statistical rejection of the model, we do not deem this finding too surprising. Indeed, the power of the typical J-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. In any case, since our model nests most of the previous estimation approaches in the literature, results remain comparable.

from what many people casually associate with the notion of peer effects, including, for example, a tendency of myopic overspending to "beat the Joneses". Of course, our results are conditional not only upon our specific interpretation of peer effects. Inevitably, they also depend upon (i) the general validity of the life-cycle framework, for which the statistical J-test provides a less comforting result than the economic "excess sensitivity" test; (ii) the specific definitions of peer groups that we consider; (iii) and our identification of peer effects based on within-peer-group differences in interest rates and demographic preference shifters, with the latter being treated as exogenous. Accordingly, it will be interesting to confront our results with complementary approaches using either an altogether different modelling context, or alternative identification approaches within the same life-cycle framework, or simply new and better consumption data as they become available.

6 Conclusion

In this paper we make two main contributions. First, we derive a suitable framework for analyzing peer effects, or "keeping up with the Joneses", in intertemporal consumption choice. 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 intertemporal consumption model. Specifically, we derive an otherwise standard Euler equation that allows individual utility to be also affected by peer-group consumption. This makes our results easily comparable to those in prominent prior studies on the dynamic patterns of household consumption, such as Attanasio et al. (1999). Focusing on the relevance of "external" habit formation, our paper also complements previous microeconomic research on "internal" habits, notably Dynan (2000).

Starting with a simple Euler equation approach, we find strong evidence for expected consumption co-movement within reference groups constructed on the basis of age, gender, race, family status, education, occupation and urbanity. This initial piece of evidence suggests an important role for consumption externalities.

However, inference under this setup is very vulnerable to misspecification, especially the omission of demographic control variables. We argue that additional restrictions are required to

discriminate between true peer-group phenomena 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.

Cast into the more robust framework of the reflection problem, our analysis suggests that much of the observed co-movement in consumption reflects correlated effects rather than true consumption externalities. In particular, the very strong peer effects suggested by the simple Euler equation approach cannot be confirmed.

Nonetheless, some of the evidence for peer effects in intertemporal consumption choice persists, as our results point to moderate consumption externalities across a range of different plausible peer-group specifications. In our view, this finding provides some comfort with respect to the many theoretical papers in modern macroeconomics that basically take peer effects for granted. Even so, more research is certainly warranted to deepen our understanding of how social interaction affects individuals' intertemporal consumption choices.

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

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

where $\text{GEOM}\left(C_t^j | X_t^j = X_t^h\right)$ 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

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

Taking logs, we obtain

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

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

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

This representation corresponds with the equation, (3), we use for modeling 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

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

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

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

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

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

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

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

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

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

with $(s^h)^2 = E_t \left[(\tilde{\varepsilon}_{t+1}^h)^2 \right]$ for all t and ζ_{t+1}^h representing a combined error term with $E_t [\zeta_{t+1}^h] = E_t \left[\left(\tilde{\varepsilon}_{t+1}^h + \frac{1}{2} (s^h)^2 - \frac{1}{2} (\tilde{\varepsilon}_{t+1}^h)^2 \right) \right] = 0$. 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.¹³ 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

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

or, rearranging terms,

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

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

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

where ϱ_{t+1}^h 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

¹³A 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.

approximation is very accurate, because all but one of the stratification variables we consider never change for a given household over time.¹⁴ 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

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

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 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

¹⁴ Age, 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.

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. The likely presence of liquidity constraints for this subpopulation might lead to a violation of the Euler equation for the respective households, thus rendering our estimation approach invalid. Deletion of the poverty subsample 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. For the estimations, we also delete the observations with the 0.75% highest and lowest consumption growth rates in order to eliminate the influence of extreme outliers. These deletions lead to a loss of an additional 396 observations. Finally, we also eliminate observations that belong to very small population strata (less than 15 observations in the baseline estimations), resulting in further data losses whose extent depends on the particular stratification at hand.

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.¹⁵ 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 (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

¹⁵See, for example, Attanasio and Browning (1995), p. 1125.

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 (5) and (8) leads to additional complications. Specifically, as reference group means $ARITM \left[\cdot | X_t^j = X_t^h \right]$ 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 important given that we would like to uncover possible peer effects at the population level rather than within our sample only.¹⁶ The next section contains a brief discussion of how we make the required adjustments.

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 $d^h \left(X_t^j \right)$ 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 $g \left(X_t^j \right)$. This step allows us to re-write the structural models (5) and (8) as

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

¹⁶Most 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.

and, respectively,

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

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 $ARITM_N \left[\cdot | g \left(X_t^j \right) \right]$ as coefficients from first-step OLS regressions of the relevant variables v , i.e. $\Delta \ln \left(C_{t+1} \right)$, ΔD_{t+1} and $\ln R_{t+1}$, on the group-year dummies $d \left(X_t^j \right)$ 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 $d \left(X_t^j \right)$. 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.¹⁷ Intuitively, sampling variability in the estimation of the reference group means is irrelevant for cases in 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.¹⁸

¹⁷In this case $F_{ARITM} = 0$ holds.

¹⁸See 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*, vol. 80(2) (May), pp. 38-42.
- Alessie, R. and Kapteyn, A. (1991). 'Habit formation, interdependent preferences and demographic effects in the Almost Ideal Demand System', *Economic Journal*, vol. 101(406) (May), pp. 404-419.
- Altug, S. and Miller, R. (1990). 'Household choices in equilibrium', *Econometrica*, vol. 58(3) (May), pp. 543-570.
- Alessie, R. and Teppa, F. (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*, vol. 1 B, pp. 741-812, Amsterdam: North Holland.
- Attanasio, O., Banks, J., Meghir, C. and Weber, G. (1999). 'Humps and bumps in lifetime consumption', *Journal of Business and Economics Statistics*, vol. 17(1) (January), pp. 22-35.
- Attanasio, O. and Browning, M. (1995). 'Consumption over the life cycle and over the business cycle', *American Economic Review*, vol. 85(5) (December), pp. 1118-1137.
- Attanasio, O. and Weber, G. (1993). 'Consumption growth, the interest rate and aggregation', *Review of Economic Studies*, vol. 60(3) (July), pp. 631-649.
- Attanasio, O. and Weber, G. (1995). 'Is consumption growth consistent with intertemporal optimization? Evidence from the Consumer Expenditure Survey', *Journal of Political Economy*, vol. 103(6) (December), pp. 1121-1157.
- Binder, M. and Pesaran, H. (2001). 'Life-cycle consumption under social interactions', *Journal of Economic Dynamics and Control*, vol. 25(1-2) (January), pp. 35-83.

- Blundell, R., Browning, M. and Meghir, C. (1994). 'Consumer demand and the life-cycle allocation of household expenditures', *Review of Economic Studies*, vol. 61(1) (January), pp. 57-80.
- Browning, M. and Collado, M. (2007). 'Habits and heterogeneity in demand: a panel data analysis', *Journal of Applied Econometrics*, vol. 22(3) (April-May), pp. 625-640.
- Browning, M. and Crossley, T. (2001). 'The life cycle model of consumption and saving', *Journal of Economic Perspectives*, vol. 15(3) (Summer), pp. 3-22.
- Campbell, J. and Cochrane, J. (1999). 'By force of habit: a consumption-based explanation of aggregate stock market behavior', *Journal of Political Economy*, vol. 107(2) (April), pp. 205-251.
- Cochrane, J. (1991). 'A simple test of consumption insurance', *Journal of Political Economy*, vol. 99(5) (October), pp. 957-976.
- Cox, D., Ng, S. and Waldkirch, A. (2004). 'Intergenerational linkages in consumption behavior', *Journal of Human Resources*, vol. 39(2) (Spring), pp. 355-381.
- Deaton, A. (1992). *Understanding Consumption*, New York: Oxford University Press.
- Duesenberry, J. (1949). *Income, Saving and the Theory of Consumer Behavior*, Cambridge, MA: Harvard University Press.
- Dynan, K. (2000). 'Habit formation in consumer preferences: evidence from panel data', *American Economic Review*, vol. 90(3) (June), pp. 391-406.
- Festinger, L. (1954). 'A theory of social comparison processes', *Human Relations*, vol. 7(2) (May), pp. 117-140.
- Gaertner, W. (1974). 'A dynamic model of interdependent consumer behavior', *Zeitschrift für Nationalökonomie*, vol. 34(3-4) (September), pp. 327-344.
- Galí, J. (1994). 'Keeping up with the Joneses: consumption externalities, portfolio choice and asset prices', *Journal of Money, Credit and Banking*, vol. 26(1) (February), pp. 1-8.

- Guariglia, A. and Rossi, M. (2002). 'Consumption, habit formation and precautionary saving: evidence from the British Household Panel Survey', *Oxford Economic Papers*, vol. 54(1) (January), pp. 1-19.
- Hall, R. and Mishkin, F. (1982). 'The sensitivity of consumption to transitory income: estimates from panel data on households', *Econometrica*, vol. 50(2) (March), pp. 461-481.
- Hayashi, F., Altonji, J. and Kotlikoff, L. (1996). 'Risk-sharing between and within families', *Econometrica*, vol. 64(2) (March), pp. 261-294.
- Heffetz, O. (2004). 'Conspicuous consumption and the visibility of consumer expenditures', Princeton University, mimeo.
- Hotz, J., Kydland, F. and Sedlacek, G. (1988). 'Intertemporal preferences and labor supply', *Econometrica*, vol. 56(2) (March), pp. 335-360.
- Hurst, E. (2004). 'Grasshoppers, ants and pre-retirement wealth: a test of permanent income consumers', University of Chicago, mimeo.
- Hurst, E. and Stafford, F. (2004). 'Home is where the equity is: liquidity constraints, refinancing and consumption', *Journal of Money, Credit and Banking*, vol. 36(6) (December), pp. 985-1014.
- Jacobs, K. (1999). 'Incomplete markets and security prices: do asset-pricing puzzles result from aggregation problems', *Journal of Finance*, vol. 54(1) (February), pp. 123-164.
- Kapteyn, A., van der Geer, S., van de Stadt, H. and Wansbeek, T. (1997). 'Interdependent preferences: an econometric analysis', *Journal of Applied Econometrics*, vol. 12 (6) (November-December), pp. 665-686.
- Lawrance, E. (1991). 'Poverty and the rate of time preference: evidence from panel data', *Journal of Political Economy*, vol. 99(1) (February), pp. 54-77.
- Leibenstein, H. (1950). 'Bandwagon, snob, and Veblen effects in the theory of consumers' demand', *Quarterly Journal of Economics*, vol. 64(2) (May), pp. 183-207.

- Ljungqvist, L. and Uhlig, H. (2000). 'Tax policy and aggregate demand management under catching up with the Joneses', *American Economic Review*, vol. 90(3) (June), pp. 356-366.
- Luttmer, E. (2006). 'Neighbors as negatives: relative earnings and well-being', *Quarterly Journal of Economics*, vol. 120(3) (June), pp. 963-1002.
- Mace, B. (1991). 'Full insurance in the presence of aggregate uncertainty', *Journal of Political Economy*, vol. 99(5) (October), pp. 928-956.
- Manski, C. (1993). 'Identification of endogenous social effects: the reflection problem', *Review of Economic Studies*, vol. 60(3) (July), pp. 531-42.
- Manski, C. (1995). *Identification Problems in the Social Sciences*, Cambridge, MA: Harvard University Press.
- Maurer, J. and Meier, A. (2005). 'Do the "Joneses" really matter? Peer-group vs. correlated effects in intertemporal consumption choice', *IFS Working Paper 15/2005*.
- Newey, W. and McFadden, D. (1994). 'Large sample estimation and hypothesis testing', in (R. Engle and D. McFadden, eds.) *Handbook of Econometrics*, vol. 4, Amsterdam: North Holland.
- Pollak, R. (1976). 'Interdependent preferences', *American Economic Review*, vol. 66(3) (June), pp. 309-320.
- Ravina, E. (2005). 'Habit persistence and keeping up with the Joneses: evidence from micro data', New York University, mimeo.
- Runkle, D. (1991). 'Liquidity constraints and the permanent-income hypothesis', *Journal of Monetary Economics*, vol. 27(1) (February), pp. 73-98.
- Shea, J. (1995). 'Union contracts and the life-cycle/permanent-income hypothesis', *American Economic Review*, vol. 85(1) (March), pp. 186-200.
- Van Praag, B. and Frijters, P. (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: Russell Sage Foundation.

Woittiez, I. and Kapteyn, A. (1998). 'Social interactions and habit formation in a model of female labour supply', *Journal of Public Economics*, vol. 70(2) (November), pp. 185-205.

Zeldes, S. (1989). 'Consumption and liquidity constraints: an empirical analysis', *Journal of Political Economy*, vol. 97(2) (April), pp. 305-346.

Discussion Paper Series

Mannheim Research Institute for the Economics of Aging Universität Mannheim

To order copies, please direct your request to the author of the title in question.

Nr.	Autoren	Titel	Jahr
155-07	Alexander Zimper, Alexander Ludwig	Attitude polarization	07
156-08	Florian Heiss, Daniel McFadden, Joachim Winter	Mind the Gap! Consumer Perceptions and Choices of Medicare Part D Prescription Drug Plans	08
157-08	Wolfgang Kuhle	Demography and Equity Premium	08
158-08	Axel Börsch-Supan, Tabea Bucher-Koenen, Anette Reil-Held, Christina Wilke	Zum künftigen Stellenwert der ersten Säule im Gesamtsystem der Alterssicherung	08
159-08	Tabea Bucher-Koenen, Christina Benita Wilke	Zur Anhebung der Altersgrenze: Eine Simulation der langfristigen Auswirkungen auf die gesetzliche Rentenversicherung bei unterschiedlichem Renteneintrittsverhalten	08
160-08	Mathias Sommer	Imputation and Harmonisation of Income, Consumption, Savings and Wealth Data from the German Income and Expenditure Survey	08
161-08	Karsten Hank	Generationenbeziehungen im alternden Europa	08
162-08	Axel Börsch-Supan, Karsten Hank, Hendrik Jürges, Mathis Schröder	Longitudinal Data Collection in Continental Europe: Experiences from the Survey of Health, Ageing and Retirement in (SHARE)	08
163-08	Martin Salm	Job loss does not cause ill health	08
164-08	Martin Salm, Daniel Schunk	The role of childhood health for the inter-generational transmission of human capital: Evidence from administrative data	08
165-08	Christina Benita Wilke	On the feasibility of notional defined contribution systems: The German case	08
166-08	Alexander Ludwig Michael Reiter	Sharing Demographic Risk – Who is Afraid of the Baby Bust?	08
167-08	Jürgen Maurer André Meier	Smooth it Like the “Joneses?” Estimating Peer-Group Effects in Intertemporal Consumption Choice	08