Review of Economic Studies (2020) **87**, 130–163 doi:10.1093/restud/rdz026

Advance access publication 6 May 2019

Consumption Network Effects

GIACOMO DE GIORGI
GSEM, University of Geneva, BGSE, BREAD, CEPR, and IPA
ANDERS FREDERIKSEN
Aarhus University, CCP, ICOA and IZA
and
LUIGI PISTAFERRI
Stanford University, NBER, CEPR, SIEPR and IZA

First version received May 2017; Editorial decision March 2019; Accepted April 2019 (*Eds.*)

In this article we study consumption network effects. Does the consumption of our peers affect our own consumption? How large is such effect? What are the economic mechanisms behind it? We use administrative panel data on Danish households to construct a measure of consumption based on tax records on income and assets. We combine tax record data with matched employer–employee data to identify peer groups based on workplace, which gives us a much tighter and credible definition of networks than used in previous literature. We use the non-overlapping network structure of one's peers group, as well as firm-level shocks, to build valid instruments for peer consumption. We estimate non-negligible and statistically significant network effects, capable of generating sizable multiplier effect at the macro-level. We also investigate what mechanisms generate such effects, distinguishing between intertemporal and intratemporal consumption effects as well as a more traditional risk sharing view.

Key words: Consumption, Networks, Social multiplier, Risk sharing.

JEL Codes: E21, D12, D85

1. INTRODUCTION

Does the consumption of our peers affect our own consumption? How large is such effect? What are the economic mechanisms behind it? What are the aggregate implications of consumption network effects? These are the questions that we investigate in this article.[1] To this purpose, we use administrative data from Denmark for the period 1980–96. The dataset includes administrative information on income and assets, so we can construct a measure of consumer spending from budget accounting. The dataset also includes information on the individual's employer ID and other observable worker characteristics, which we use to construct reference groups made of co-workers sharing similar characteristics (such as occupation or education). Finally, we can match our administrative dataset with a small consumption survey where we observe household expenditures on various goods. This allows us to distinguish between competing mechanisms regarding the economic interpretation of consumption network effects.

1. We will use the terms "peer effect" and "network effect" interchangeably, although the latter is better used in a context in which the utility from consuming a certain good is a function of the number of consumers (either because of congestion or economies of scale).

The editor in charge of this paper was Aureo de Paula.

The study of social influences on consumption behaviour has a long history in economics, dating back at least to Veblen (1899). Duesenberry (1948) also emphasized the role of social influences on consumption in his relative income hypothesis. In recent years, the study of social influences on individual behaviour has grown substantially to include effects on education, welfare participation, job search, work effort, and workers' productivity among other things (see Bertrand *et al.*, 2000; Duflo and Saez, 2003; Mas and Moretti, 2009; De Giorgi *et al.*, 2010, for recent contributions). In macroeconomics, variants of the "keeping-up-with-the-Joneses" idea have been used to explain portfolio choice, growth, and tax policy among other things (see *e.g.* Gali, 1994; Carroll *et al.*, 1997; Ljungqvist and Uhlig, 2000).

The study of social influences on consumption behaviour has evolved along two different lines. First, the definition of the relevant reference group. Here, empirical work has been mostly constrained by the type of consumption data available (typically, small consumption surveys with little or no longitudinal component). Hence, peers have been defined generically as individuals sharing similar socio-demographic characteristics (as in Maurer and Meier, 2008), or somewhat more precisely as a racial group within a U.S. state (Charles *et al.*, 2009) or narrower geographical entities (Kuhn *et al.*, 2011; Agarwal *et al.*, 2017). Second, the literature has proposed several economic explanations for the underlying estimated peer effects. In models of consumption behaviour, agents decide the allocation of goods within the period as well as over the life cycle. It becomes then natural to ask whether the influence of peers is on the demand for specific goods (*i.e.* through intratemporal effects), or on altering the amount of overall consumption allocated to different periods (*i.e.* through an intertemporal effect). Alternatively, the theory of risk sharing suggests that risks are shared among members of a reference group, creating correlation among their consumptions.

Our article advances and contributes to both lines of research. First, we assume that co-workers are the relevant reference group of individuals and reconstruct the social network of a given household using information on the husband's and the wife's workplace. In the empirical analysis, we define as co-worker someone who works in the same plant and is "similar" in terms of occupation and education. Co-workers represent a naturally occurring peer group. Indeed, co-workers tend to spend a substantial fraction of their time together. Moreover, friendship often causes co-workership due to job search strategies adopted by job seekers (Montgomery, 1991).

Our second contribution is to propose and implement empirical tests that allow us to study whether peer effects emerge because of intertemporal substitution in consumption across periods, substitution across goods within a given period, or because of risk sharing among members of a network.

Why is the study of consumption network effects important? There are at least two reasons. First, from a welfare point of view one may be interested in measuring and understanding the type of distortions (if any) induced by the presence of peer effects. Depending on the mechanism underlying peer effects, distortions may be intratemporal and/or intertemporal. In the first case, budget shares would be distorted. For example, status-seeking behaviour (as in conspicuous consumption models) might inflate the share of conspicuous goods over the consumption bundle. Since conspicuous goods are typically luxuries (cars or jewellery being the most notable examples), consumption peer effects might have noticeable welfare consequences (in the form of excess "wasteful" consumption). In the intertemporal case, the saving profile would be different from the optimal one we would observe when agents act atomistically. This may induce undersaving (or over-borrowing) in the attempt to keeping up with the peers. Finally, if risk sharing is the main reason for correlated consumption profiles, we would actually record important welfare gains.

The second reason why studying consumption network effects is important is because of their potential aggregate effects. Uninsured idiosyncratic shocks (such as unanticipated tax changes targeting rich taxpayers) might have aggregate consequences that go beyond the group directly affected by the shock. This depends on the size of the estimated effect as well as the degree of connectedness between groups that are directly affected and unaffected by the shock. In our empirical analysis, we find non-negligible endogenous peer effects, which translate into a non-negligible social multiplier. We then analyse the effect of policy counterfactuals based on hypothetical consumption stimulus programs targeting different groups in the population.

While the economic issues regarding the presence and importance of consumption peer effects are not trivial (as they may be consistent with different theoretical mechanisms), the econometric issues surrounding identification of such effects are no less trivial, as is well known at least since Manski (1993). In particular, identification of consumption peer effects in a linear-in-means model is difficult because peers may have similar levels of consumption due to: (a) contextual effects, (b) endogenous effects, or (c) correlated effects. In our specific application, these three effects could be described as follows: (a) workers with highly educated peers may have different wealth accumulation attitudes than those with mostly low-educated peers; (b) there may be genuine peer influences, *i.e.*, consumption behaviour changes (causally) in response to the consumption behaviour of co-workers; and finally, (c) consumption of all workers within the firm may be affected by some common (firm-level) unobserved shock, such as a firm-level productivity change or a health campaign within the firm. In principle, one can estimate effect (a) using random assignment as in Sacerdote (2001) or De Giorgi *et al.* (2010). However, random assignment does not alone allow consistent estimation of effects (b) or (c).

We tackle these econometric issues by extending the network approach idea of Bramoullé *et al.* (2009) and De Giorgi *et al.* (2010) with the use of exogenous shocks to distance-3 nodes. This requires the existence of intransitive triads, *i.e.*, "friends of friends who are not friends themselves". However, since this idea is often opaque in its practical implementation, we justify it economically with the use of firm-level idiosyncratic variation. To give a simple example, our identification strategy rests on the idea that an event like a firm downsizing experienced by the co-worker of the spouse of my co-worker (controlling for common shocks experienced by my firm) has no direct effect on my consumption but only an indirect one (through peer effects).

In our specific context, the key (and novel) fact that we exploit empirically is that working relationships are individual, but consumption is shared among spouses. Hence, spouses add nodes to otherwise unconnected networks (*i.e.* groups of workers sharing similar characteristics within a firm). It follows that exogenous variation affecting the consumption of the co-workers of the spouses of husband's and wife's co-workers represent valid exclusion restrictions.

Our IV strategy delivers an estimate of the elasticity of own consumption with respect to peers' consumption of about 0.3, which is statistically indistinguishable between husband's and wife's.[2] Such an estimated effect translates into a non-trivial aggregate effect which depends upon the degree of connectedness of the households. When we explore the theoretical mechanism behind our results, we find support for peer effects acting primarily on the intertemporal aspect of consumption choice, while we can rule out sharp versions of models with intratemporal distortions as well as full and partial risk sharing (with the caveat that our test for intratemporal effects has lower power than the test for intertemporal effects).

The rest of the paper is organized as follows. In section 2 we provide information on the data we have available. In section 3 we consider three different economic mechanisms that may potentially generate a relationship between individual consumption and the consumption of peers, and discuss testing strategies that allow us to distinguish between them. Section 4 is devoted to a discussion of the identification strategy and Section 5 to the results. Section 6 discusses the results of a simple simulation of the aggregate implications of our findings, while section 7 concludes.

2. The response to a random peer's consumption is much smaller due to large network size.

2. DATA

2.1. *Tax records data matched with employee records*

We use administrative longitudinal tax records for the Danish population for the 1980–96 period. Chetty *et al.* (2011) provide an informed discussion of the Danish tax system. The dataset includes information on income and assets for each taxpayer. During this period information on income and assets (with the exception of durables such as cars, jewellery, etc.) come from third-party reports (*e.g.* from employers, banks, stockbrokers, etc.), thus minimizing measurement error. While income data are typically available in all tax record datasets, asset data are available because, until 1996, households were subject to a wealth tax.[3] We match the administrative data with the Integrated Database for Labor Market Research (IDA), an employer–employee dataset, which includes, among other things, demographics and firm and plant ID's, from which we can identify co-workers. We define co-workers as individuals who work in the same plant (for public employees, this is the physical address of their workplace)—see below for more precise definition.

Our estimation sample includes households whose head is aged 18–65, where both spouses work and are employees rather than self-employed. We no longer use these households if one or both members stop working or become self-employed.[4] Table A1 in the online Appendix details the step-by-step selection process. Most of the observations are lost when focusing on working individuals, those who are married, and those whose spouse also works. Our sample selection strategy is driven by the research objective—we can only identify the reference network if people are employed; and we can only form instruments if spouses also work. However, we stress that in the computation of peers' consumption (the right hand side variable of our regressions) we use *all* workers, including singles and households with only one spouse working.

Consumption is not directly measured in administrative tax data. We use the dynamic budget constraint to calculate total consumption (or more precisely, total spending). In particular, consumption is calculated as the difference between after-tax annual income and asset changes:

$$C_{it} = Y_{it} - T_{it} - \Delta A_{it}, \tag{1}$$

where $Y_{it} = (GY_{it} + HS_{it} + CS_{it} - TH_{it})$, GY is gross income (the sum of income from all sources, labour and capital), HS the value of housing support, CS the value of child support, TH the implicit tax on the consumption value of owned housing, T the total tax payments, and ΔA the change in asset values (defined as the sum of cash, deposits on bank accounts, stocks and shares, the value of property, and the value of cars and other types of vehicles minus liabilities). This is similar to Browning and Leth-Petersen (2003), who conclude that this simple measure tends to behave as well as more sophisticated measures which attempt to account for capital gains, etc. (see below for a formal comparison with survey data). Note that this measure is robust to cases in which consumers enjoy different returns on their financial investments, as Y_{it} includes capital income, which incorporates directly such return heterogeneity (see Fagereng *et al.*, 2016). In some of the robustness exercises below, we investigate the sensitivity of the results to dropping households for whom capital gains or losses may be important, such as stockholders or homeowners.

Table 1 provides some descriptive statistics about our sample. All monetary values are in 2000 DKr prices (for comparison, 1\$=7.33DKr). Annual (before-tax) household income (Y_{it} above) is

3. Tax record data are actually available for later years, but the wealth tax was abolished in 1996. Collection of detailed asset data was thus discontinued after 1996. In particular, after 1996 only third-party reports remain available to researchers.

4. Given the applied selection criteria, less than 5% of the households ever change spousal composition in our sample period. Hence, we abstract from divorce and separation in the analysis.

TABLE 1
Descriptive statistics

	Mean	Std. Dev.		Mean (Median)	Std. Dev.
Outcomes					
log Consumption	12.07	0.66	Consumption	358,873 (314,727)	324,117
log Income	12.57	0.32	Income	515,877 (484,436)	186,305
log Disposable income	12.16	0.28	Disposable income	340,041 (325,877)	117,353
			Assets	226,567 (141,792)	758,139
Socio-demographics:					
Age			Sector: Manufacturing		
Husband	42.53	9.42	Husband	25.14	
Wife	40.06	9.10	Wife	12.75	
Years of schooling			Sector: Service		
Husband	12.06	2.33	Husband	15.63	
Wife	11.70	2.33	Wife	12.22	
Occupation: Blue			Sector: Construction		
Husband	43.04		Husband	10.30	
Wife	31.63		Wife	0.99	
Occupation: White			Sector: Other		
Husband	15.83		Husband	48.93	
Wife	45.20		Wife	74.05	
Occupation: Manager			Tenure (in 1996):		
Husband	41.14		Husband	4.79	4.92
Wife	23.18		Wife	4.68	4.94
# Kids 0–6	0.38	0.66	# Kids 7–18	0.72	0.86
Workplace characteristics:					
Firm size			Type: Publicly traded		
Husband	260	65	Husband	0.46	
Wife	330	82	Wife	0.24	
Change in firm size			Type: Limited liability		
Husband	−9	32	Husband	0.08	
Wife	−13	42	Wife	0.04	
Public sector			Type: Other		
Husband	0.32		Husband	0.46	
Wife	0.61		Wife	0.72	
Number of households: 757,439					

about DKr 516k; after-tax income is DKr 340k. The value of assets (about DKr 227k) is smaller than what would typically be recorded in the U.S., although we note that there is quite a large dispersion in asset values (a standard deviation of about DKr 760k). Consumption is about DKr 359K. Note that the average of our consumption measure is influenced by a long right tail, as evidenced by the large discrepancy between mean and median values and the large standard deviation (relative to before- or after-tax income), which partly derives from the imputed nature of the variable and partly from the fact that spending is more volatile than consumption. In terms of socio-demographic characteristics, husbands are slightly older than wives and slightly more educated. There is a large concentration of women in "white collar" jobs, and a larger concentration of men in "managerial" and "blue-collar" positions relative to females. As for sectorial concentration, there is a higher proportion of men in manufacturing and construction, and a higher proportion of women in services and "other sectors" (mostly, public employment). We also compute tenure (years with current employer within our sample period 1980–96) and do not find large differences across genders (5 years on average). This tells us that co-workers tend to be in the same firm/location for a non-negligible number of years. Finally, the households in our sample have on average 0.4 young children (0–6 years of age) and 0.7 older children (7–18 years old).

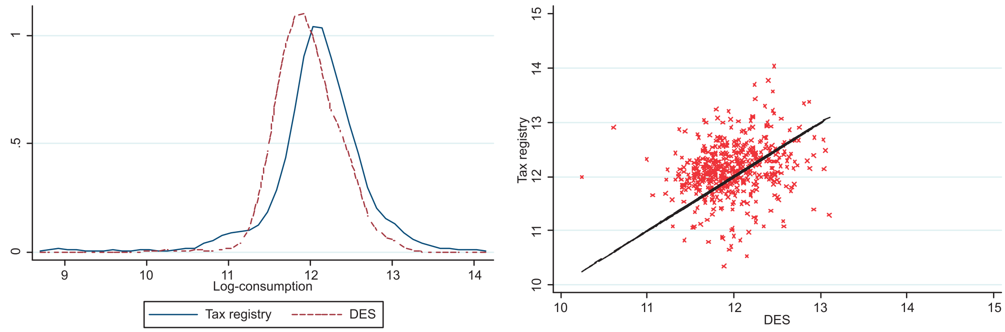

FIGURE 1

The distribution of consumption in the Tax Registry and in the Danish Expenditure Survey.

At the bottom of Table 1, we also report workplace characteristics that we use as controls and instruments. Average firm size is 260 and 330 for husbands and wives, respectively. The annual employment change is negative in this period, although the standard deviation is rather high. As mentioned above, a larger fraction of women work in the public sector (60%) relative to males (32%). Men tend to be more represented in publicly traded company than women (46% versus 24%). A similar pattern emerges for limited liability companies with a larger share of men (8% versus 4%), while the pattern is reversed in "other companies" (mostly located in the public sector), where the fraction of women and men working are 72% and 46%, respectively.

2.2. *Danish Expenditure Survey*

The Danish Expenditure Survey (DES) is, in (relative) size and scope, very similar to the U.S. Consumer Expenditure Survey (CEX) or the U.K. Family Expenditure Survey (FES). We use the DES mostly to investigate the economic mechanisms behind our findings. See Browning and Leth-Petersen (2003) for more details about the survey. The survey is available from 1994, but given that our administrative data end in 1996, we use only the three waves spanning 1994–1996 (note that the spending data are not longitudinal), and focus on the households that can be matched with the tax records.

The left panel of Figure 1 plots the consumption density in the Tax Registry and in the survey data (both expressed in logs). The two distributions overlap significantly and differ appreciably only in the tails (due to issues related to capital gains and losses that are hard to account for in the Tax Registry data). In one of the robustness exercises below, we investigate the sensitivity of the results to removing the tails of the consumption distribution or focus on samples where capital gains and losses are unlikely to be important. In the right panel of Figure 1 we plot, for the households observed in both the tax records and the DES, the tax-record consumption measure against the survey consumption measure, along with the 45° line. As in the distributional plot, the tax-record measure appears to have a longer tail.[5]

To conduct the tests we describe in the next section, we divide spending in the DES into spending on visible, neutral, and not-visible goods. While for most goods the separation is unambiguous (*i.e.* jewellery or home insurance), we use an index of visibility proposed by

5. Since the two variables can be interpreted as two error-ridden measures of true consumption, there is no reason to expect a one-to-one relationship. The measurement errors are of different nature (in the tax records, they come from the inability to measure capital gains and losses, etc.; in the survey, from traditional recall bias, understatement of certain expenditures, etc.). In unreported regressions, we verify that the difference between the two consumption measures cannot be explained by observable characteristics, implying that none of the difference appears "systematic".

Heffetz (2011) as an anchor. To construct the index, Heffetz (2011) conducts an original survey where each respondent is asked to rank 31 categories of expenditure according to their external "visibility". The higher the visibility, the higher the assumed conspicuousness. We define visible goods to include Tobacco and Alcohol, Food away from home, Clothing, Furniture and Home goods, Electrics/Appliances, Vehicles, Entertainment, Books, Education, Personal care. Neutral is limited to food at home. Everything else is classified as non-visible (insurance, rent, etc.). In an extension of the testing idea, we construct spending categories that reproduce exactly the separation proposed by Heffetz (2011), with the exception of charity contributions that are not observed in the DES. See Table A6 in the Appendix for more details.

3. GENERAL THEORETICAL FRAMEWORK

In this section, we explore the theoretical mechanisms that may be responsible for the presence of consumption network effects. In general, one can think of network effects inducing either shifts in individual preferences or shifts in individual resources. In this section, we discuss the first type of effects, and in Section 3.2 we discuss the second type of effects.

3.1. *Intratemporal versus intertemporal distortions*

To formally analyse network effects in a traditional life cycle consumption framework, we assume that the problem of the consumer can be written as:

$$\max E_0 \sum_{t=0}^{T} U_t(\mathbf{p}_t, C_{it}, z_{it})$$

subject to the dynamic budget constraint:

$$A_{it+1} = (1+r)(A_{it} + Y_{it} - C_{it}),$$

where $C_{it} = \sum_{k=1}^{K} p_t^k q_{it}^k$ is total spending on goods q_{it}^k with prices p_t^k $(k = 1 \ldots K)$, A_{it} is assets, Y_{it} income, and r the interest rate. E_0 represents expectations conditional on information available at time 0.

We follow Blundell *et al.* (1994) in considering a general form for the conditional indirect utility function $U_t(.)$:

$$U_t(\mathbf{p}_t, C_{it}, z_{it}) = F_t\left(V_t\left(\mathbf{p}_t, C_{it}, z_{it}^1\right), z_{it}^2\right) + G\left(z_{it}^3\right). \tag{2}$$

In this setting $V_t(.)$ governs the within-period allocation of total spending C_{it} to goods q_{it}^k, while U_t determines the intertemporal (or between-periods) allocation (*i.e.* the choice between consumption and savings). $F_t(.)$ is a strictly increasing monotonic transformation. Finally, $z_{it} = (z_{1it}, z_{2it}, z_{3it})$ is a vector of conditioning goods or characteristics (with z_{1it}, z_{2it}, and z_{3it} possibly having overlapping terms). We can think of peers' consumption (or the composition thereof) as being one such conditioning characteristic. The term z_{it} would typically include observed or unobserved taste shifters and, in some of the contexts studied in the literature, labour supply or demographics (see *e.g.* Blundell *et al.*, 1994). In what follows, we assume:

$$z_{it} = \left(X_{it}, \zeta_{it}, \{C_{nt}\}_{n=1, n\neq i}^{N}\right),$$

where X_{it} and ζ_{it} are observed and unobserved taste shifters, respectively, and $\{C_{nt}\}$ is the vector of peers' consumption.[6]

In principle, peers' consumption can enter any aspect of the consumption problem. To look at cases of interest, we start by noting that the demand functions (representing *intratemporal* or within-period allocation) are independent of $F_t(.)$ and are hence determined by the usual Roy's identity:

$$q_{it}^k = -\frac{\frac{\partial V_t(.)}{\partial p_t^k}}{\frac{\partial V_t(.)}{\partial C_{it}}}. \tag{3}$$

In contrast, the Euler equation (representing *intertemporal* or between-period allocation) is given by:

$$E_t \frac{\partial U_{t+1}(.)}{\partial C_{it+1}} = (1+r)^{-1} \frac{\partial U_t(.)}{\partial C_{it}} \tag{4}$$

or $E_t \frac{\partial F_{t+1}}{\partial V_{t+1}} \frac{\partial V_{t+1}}{\partial C_{it+1}} = (1+r)^{-1} \frac{\partial F_t}{\partial V_t} \frac{\partial V_t}{\partial C_{it}}$. We can now consider three cases of interest.

CASE 1: Additive separability:

$$U_t(\mathbf{p}_t, C_{it}, z_{it}) = F_t(V_t(\mathbf{p}_t, C_{it}, X_{it}, \zeta_{it})) + G\left(\{C_{nt}\}_{n=1, n\neq i}^N\right)$$

In this case, application of (3) yields $\frac{\partial q_{it}^k}{\partial C_{nt}} = 0$, since $V_t(.)$ does not depend on C_{nt} for all $n \neq i$ and all $k = \{1, 2, \ldots, K\}$. Hence the intratemporal allocation is independent of peers' consumption. Since $\frac{\partial V_s}{\partial C_{ns}} = 0$ for all s, the intertemporal allocation decision is also independent of peers' consumption. Therefore, in this case, there are no network effects on consumption.

CASE 2: Weak intratemporal separability:

$$U_t(\mathbf{p}_t, C_{it}, z_{it}) = F_t\left(V_t(\mathbf{p}_t, C_{it}, X_{it}, \zeta_{it}), \{C_{nt}\}_{n=1, n\neq i}^N\right)$$

As before, $\frac{\partial q_{it}^k}{\partial C_{nt}} = 0$ because $V_t(.)$ does not include C_{nt} (for all $n \neq i$). Hence intratemporal allocation is again independent of peer consumption when C_{nt} enters preferences as weakly separable, *as long as* one conditions on within-period spending C_{it}. This is a powerful testable restriction, similar in spirit to the one proposed by Browning and Meghir (1991) in a different context.

In contrast, the marginal utility of total consumption changes with peers' consumption, inducing *intertemporal* effects. To see this with a concrete example, consider a simple functional form (similar to the one proposed by Blundell *et al.*, 1994):[7]

$$U_t(.) = (1+\phi)^{-t} \frac{(C_{it}/a(\mathbf{p}_t))^{1-\rho} - 1}{1-\rho} \frac{\exp\{\kappa X_{it} + \zeta_{it}\}}{b(\mathbf{p}_t)} \prod_{n=1, n\neq i}^{N} C_{nt}^{\lambda}.$$

6. Note that z_{it} may also include the vector of peer's exogenous characteristics $\{X_{nt}\}_{n=1, n\neq i}^N$, which we omit for simplicity. The idea is that being surrounded by many educated people (say) may shift the preferences for consumption versus savings, or preferences for some specific goods.

7. The functions $a(\mathbf{p})$ and $b(\mathbf{p})$ are linear, positive, and homogeneous. They can be interpreted as the costs of subsistence and bliss, respectively. See Deaton and Muellbauer (1980).

The (log-linearized) Euler equation is (approximately):

$$\Delta \ln \frac{C_{it+1}}{a\left(\mathbf{p}_{t+1}\right)} \cong \rho^{-1}\left((r-\phi)+\kappa \Delta X_{it}-\Delta \ln b\left(\mathbf{p}_{t+1}\right)+\lambda \Delta \frac{\overline{\ln C_{it+1}}}{a\left(\mathbf{p}_{t+1}\right)}\right)+\varepsilon_{it+1}, \tag{5}$$

where $\frac{C_{it+1}}{a(\mathbf{p}_{t+1})}$ is real consumption expenditure and ε_{it+1} includes both unobserved taste shifters (the ζ_{it} term defined above), as well as unanticipated changes in household resources (income shocks). In our empirical framework below, some of these shocks will be related to firm variables. There is a growing literature documenting the importance of firm effects for wages and changes thereof (see *e.g.* Guiso *et al.*, 2004; Card *et al.*, 2014; Kline *et al.*, 2017). If firm shocks pass through the permanent component of the worker's wage (perhaps due to labour market frictions preventing workers' wages to reflect their marginal productivity), they will naturally pass through consumption as well.

Equation (5) shows that consumption allocation across periods depends on peers' consumption (as long as $\lambda \neq 0$). Hence, an increase in peer consumption may change the allocation between consumption and savings (induce under- or over-saving) relative to the case $\lambda = 0$.

CASE 3: Intratemporal non-separability:

$$U_t(\mathbf{p}_t, C_{it}, z_{it}) = F_t\left(V_t\left(\mathbf{p}_t, C_t, X_{it}, \zeta_{it}, \{C_{nt}\}_{n=1,n\neq i}^{N}\right)\right).$$

Assume for example that:

$$V_t(.) = (1+\phi)^{-t} \frac{\left(C_{it}/a\left(\mathbf{p}_t, \{C_{nt}\}_{n=1,n\neq i}^{N}\right)\right)^{1-\rho}-1}{1-\rho} \frac{\exp\{\kappa X_{it}+\zeta_{it}\}}{b\left(\mathbf{p}_t, \{C_{nt}\}_{n=1,n\neq i}^{N}\right)} \prod_{n=1,n\neq i}^{N} C_{nt}^{\lambda}$$

From now on, we denote: $a_t(.) = a\left(\mathbf{p}_t, \{C_{nt}\}_{n=1,n\neq i}^{N}\right)$ and $b_t(.) = b\left(\mathbf{p}_t, \{C_{nt}\}_{n=1,n\neq i}^{N}\right)$ to avoid cluttering. In this third case, application of Roy's identity gives the budget share on good j:

$$\omega_{it}^{j} = \frac{p_t^j q_{it}^j}{C_{it}} = \frac{\partial \ln b_t(.)}{\partial \ln p_t^j} \frac{1-(C_{it}/a_t(.))^{-(1-\rho)}}{1-\rho} + \frac{\partial \ln a_t(.)}{\partial \ln p_t^j}.$$

Intratemporal allocations will now be distorted by peers' consumption if the latter shifts the price elasticity of goods. For example, if we adopt a simple linear shifter specification:

$$\ln a_t(.) = \alpha_0 + \sum_k \left(\alpha_{0k} + \alpha_{1k}\overline{\ln C}_{it}\right)\ln p_t^k + \frac{1}{2}\sum_k \sum_j \eta_{kj} \ln p_t^k \ln p_t^j$$

$$\ln b_t(.) = \sum_k \left(\beta_{0k} + \beta_{1k}\overline{\ln C}_{it}\right)\ln p_t^k,$$

then spending on good j will depend on peers' consumption according to the sign and magnitude of the coefficients α_{1j} and β_{1j}. For example, with the functional form above, the budget share for good j is:

$$\omega_{it}^{j} = \alpha_{0j} + \alpha_{1j}\overline{\ln C}_{it} + \sum_k \eta_{jk} \ln p_t^k + \left(\beta_{0j} + \beta_{1j}\overline{\ln C}_{it}\right) \frac{1-(C_{it}/a_t(.))^{-(1-\rho)}}{1-\rho} \tag{6}$$

TABLE 2
Does $\overline{\ln C}$ enters the demand functions or the Euler equation?

	$\overline{\ln C} \in z^3, \overline{\ln C} \notin \{z^1, z^2\}$	$\overline{\ln C} \in z^2, \overline{\ln C} \notin z^1$	$\overline{\ln C} \in z^1, \overline{\ln C} \notin z^2$
Demand functions	No	No	Yes
Euler equation	No	Yes	Yes

As for intertemporal allocation, they are also distorted, as the Euler equation is now:

$$\Delta \ln \frac{C_{it+1}}{a_{t+1}(.)} \cong \rho^{-1}\left((r-\phi)+\kappa\,\Delta X_{it} - \Delta \ln b_{t+1}(.) + \lambda \Delta \frac{\overline{\ln C_{it+1}}}{a_{t+1}(.)}\right) + \varepsilon_{it+1} \tag{7}$$

where, as before, ε_{t+1} includes both unobserved taste shifters and income shocks.[8]

From the general form $U_t(\mathbf{p}_t, C_t, z_t) = F_t\left(V_t\left(\mathbf{p}_t, C_t, z_t^1\right), z_t^2\right) + G\left(z_t^3\right)$, Table 2 summarizes the possible cases we can confront. Our strategy for distinguishing between these various cases is sequential. First, we estimate Euler equations for individual consumption growth that control for peers' consumption growth. Given that we do not observe good-specific prices, we will proxy the indexes $a_{t+1}(.)$ and $b_{t+1}(.)$ with a full set of year dummies and region dummies. If we find no peer effects, we can conclude that preferences are intratemporally additive separable. If we find that peer effects are present (which as we shall see is the relevant empirical case), we need to distinguish between the case in which distortions are only intertemporal, or the case in which distortions are both inter- and intra-temporal.

We can distinguish between these two cases by estimating demand functions and testing whether peers' consumption can be excluded from the demand for the various goods considered (controlling, crucially, for private total spending). In other words, we can estimate (6) and test whether $\alpha_{1j} = 0$ and $\beta_{1j} = 0$. Since the most prominent theory for justifying the presence of intratemporal distortions is the "conspicuous consumption" hypothesis, we divide goods according to their degree of conspicuousness (*i.e.* "visible" versus "less visible" goods). An additional implication of the conspicuous consumption hypothesis (discussed above) is that we should observe "reshuffling" (see Section A.2 in the Appendix). If we reject both the presence of peers' consumption and reshuffling, then we can conclude that distortions are only intertemporal.

The estimation strategy assumes that we can obtain consistent estimates of consumption peer effects. This is notoriously difficult due to a host of identification problems remarked in the peer effects literature. We discuss how the structure of networks (at the co-worker level), as well as the use of exogenous firm-level shocks, helps us achieve identification in the next section. Once we have established what the main theoretical mechanism is (if any), we investigate its magnitude, heterogeneity, and robustness. Finally, we discuss welfare and macroeconomic implications.

3.2. *Risk sharing*

A final theory for why consumptions can be correlated across agents is because of risk sharing among co-workers. Workers' repeated interactions in the workplace may indeed favour risk pooling. In full insurance versions of the theory, the growth rates of consumption of people belonging to the same risk sharing pool are perfectly correlated (Mace, 1991). Hence, full

8. In models with "conspicuousness" researchers draw a difference between "visible" and "non-visible" goods. In the Appendix, we use a simplified functional forms for $a_t(.)$ and $b_t(.)$ to show that this induces "reshuffling behaviour" as peers' consumption increases: the demand for visible goods increases and that for goods that are not visible declines.

insurance implies $\rho^{-1}\lambda = 1$ when estimating an equation like (7). Note that in this case there is no meaningful "causal" relationship running from consumption of peers to individual consumption. The levels of consumption of individuals sharing risks optimally grow at the same rate because the effect of idiosyncratic shocks has been neutralized.

However, full insurance is an extreme view of risk sharing, especially in a setting like ours in which there is substantial social insurance provided by the Danish welfare system. It is more likely that, if risk sharing among co-workers exists, it provides only partial insurance.

Suppose that risk sharing is implemented *via* side transfers, *i.e.*, workers receive transfer payments in bad times while the flow is reversed in good times. If worker i has been unlucky ($\Delta \ln Y_{it} < 0$) and co-worker j has been lucky ($\Delta \ln Y_{jt} \geq 0$), worker j would transfer to i some payments that go unrecorded in the tax record definition of consumption. This means that consumption in the tax records (C^T) systematically understates true consumption for the unlucky workers and systematically overstates it for the lucky workers. However, the consumption definition coming from the consumer survey (C^S) will fully reflect transfers because it is based on actual spending on goods (which is partly financed by transfers received or paid). It follows that the difference $(\ln C^S_{it} - \ln C^T_{it})$ will be systematically *negatively* correlated with $\Delta \ln Y_{it}$ (controlling for $\Delta \ln Y_{jt}$ or, in line with our application, for $\Delta\overline{\ln Y_{it}}$) if risk sharing considerations are at play. Similarly, $(\ln C^S_{it} - \ln C^T_{it})$ will be systematically *positively* correlated with $\Delta \ln Y_{jt}$ ($\Delta\overline{\ln Y_{it}}$ in the empirical test) once we control for individual income growth $\Delta \ln Y_{it}$. Hence, we can run a regression:

$$\ln C^S_{it} - \ln C^T_{it} = \pi_0 + \pi_1 \Delta \ln Y_{it} + \pi_2 \Delta\overline{\ln Y_{it}} + \upsilon_{it},$$

(where υ_{it} is an error term reflecting the differences with which consumption data are measured in the two samples) and test whether $\pi_1 < 0$ and $\pi_2 > 0$.[9]

4. IDENTIFICATION

Identifying consumption network or social interaction effects is not trivial. Two problems in particular need to be tackled. First, the definition of the relevant network or reference group. Second, the endogeneity of the peers' consumption variable.

The definition of networks or reference groups in economics is difficult and severely limited by data availability (see De Paula, 2017). Ideally, one would survey individuals, reconstruct the web of interactions they span (family, friends, co-workers, etc.), and then collect socio-economic information on both ends of each edge. In practice, this is a rarely accomplished task (exceptions are the Add Health data in the U.S.; and the Indian microfinance clients network of Banerjee *et al.* (2013)), and identification of networks proceeds instead with identifying characteristics that are common to all network members (such as race, neighbourhood, classroom, cohort, and interactions thereof).[10] In this article, we assume that individuals who work together form a social network. There are two reasons why co-workers may be a more credible reference group than the definitions adopted in the consumption literature. First, if social effects increase with the time spent with members of the reference group, "co-workers" are obvious candidates for the

9. Note that the test that $\pi_2 > 0$ may be more robust than the test that $\pi_1 < 0$. The reason is that there may be a spuriously negative correlation between $\left(\ln C^S_{it} - \ln C^T_{it}\right)$ and $\Delta \ln Y_{it}$. Suppose that $\ln C^T_{it}$ includes spending on durables or capital gain and $\ln C^S_{it}$ does not. When $\Delta \ln Y_{it}$ grows, people may buy more durables, which may induce a negative correlation between $\left(\ln C^S_{it} - \ln C^T_{it}\right)$ and $\Delta \ln Y_{it}$ that is unrelated to risk sharing considerations.

10. More recently, researchers have tried to reconstruct the strength of social connections directly from outcome data (De Paula *et al.*, 2018).

ideal reference group, as they are the individuals people spend most of their day with. Second, in principle the ideal peer is a "friend". Evidence from sociology and labour economics shows that finding jobs through friends is one of the most frequent job search mechanisms utilized by job-seeker workers (Holzer, 1988). Hence, not only do co-workers become friends; in some cases it is actually friendship that causes co-workership. Nonetheless, our definition of network may identify the true network of an individual only imperfectly: some co-workers do not exert any social influence, and other non-co-workers may play an important social role. Hence, our networks can be measured with error. The IV strategy we define below is designed to correct for this problem, as well as for the measurement error in our consumption measure, on top of the standard endogeneity problems.

Identification of peer effects is plagued by a number of econometric issues (Manski, 1993; Brock and Durlauf, 2001; Moffitt, 2001). In general, the popular linear-in-means model allows for three distinct type of effects: (a) contextual effects, (b) endogeneous effects, and (c) correlated effects. Contextual effects may emerge if co-workers share traits that make them more likely to select a given firm and these traits are important determinants of the dependent variable under study. Endogenous effects are the genuine network effects we are interested in. Finally, correlated effects may emerge if workers share unobserved shocks (say, a cut in their wages due to a firm productivity shock) that make their consumption move simultaneously *independently* of any genuine network effects. In general, when all effects are present it is very hard to distinguish one's behaviour as cause or effect of someone else's behaviour. In the same vein if similar individuals or households have common behaviour, it is very hard to say whether this is because they are very similar to start with or because they are influencing each other.

Our identification strategy relies on exploiting the social network structure of the households in our sample. The main idea is that individuals are part of social networks that overlap only imperfectly (as in Bramoullé *et al.*, 2009; Calvó-Armengol *et al.*, 2009; De Giorgi *et al.*, 2010). In our specific context, we use the fact that social relationships are established along two lines: at the family level (*e.g.* husband and wife) and at the firm level (co-workers). If husband and wife work in different firms, it is possible to construct intransitive triads, *i.e.*, "friends of friends who are not friends themselves". As we shall illustrate in what follows, this allows identification of all parameters of interest of the model.

More formally, we consider the following linear-in-mean specification for consumption growth, which is a simple generalization of the Euler equation (7) above (to allow for multiplexity, *i.e.*, the fact that—at least in principle—husband and wife can have distinct networks):

$$\begin{aligned}\Delta \ln C_{it} = \delta_0 + \theta_w \Delta \overline{\ln C_{it}^w} + \theta_h \Delta \overline{\ln C_{it}^h} + \gamma_w \Delta \overline{X_{it}^w} + \gamma_h \Delta \overline{X_{it}^h} \\ + \delta_w \Delta X_{it}^w + \delta_h \Delta X_{it}^h + \varepsilon_{it}.\end{aligned} \tag{8}$$

Here i and t indicate household and time, while the superscripts w and h indicate wife and husband, respectively. Hence, $\overline{\ln C_{it}^w}$ and $\overline{\ln C_{it}^h}$ are the (average) log consumption levels of the wife's and husband's co-workers; $\overline{X_{it}^w}$, $\overline{X_{it}^h}$ are the (average) characteristics of the wife's and husband's co-workers which can also include firm-level variables, such as a firm size or other characteristics; X_{it}^w, X_{it}^h are the wife's and husband observable characteristics. In keeping with the specification in equation (7), all variables are expressed in first differences. There are a series of good reasons why one might want to consider the two spouses' networks separately, *e.g.*, differential preferences, differential strength of social influence by gender, as well as different bargaining power within the household. We will not make any attempt to micro-found our analysis as the bulk of our data comes from the administrative tax records, and therefore we only measure

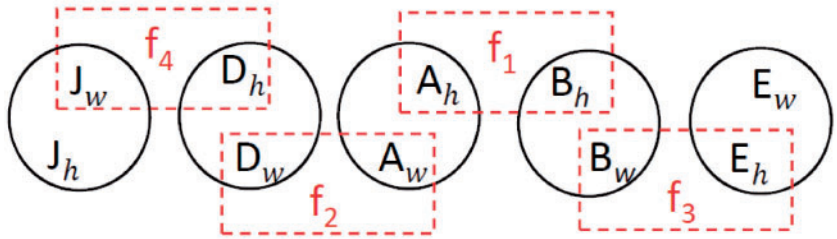

FIGURE 2
A simple example of network identification.

total expenditure at the household level (see equation 1). Moreover, we lack information on labour supply.[11]

The main parameters of interest in (8) are the θ's (endogenous effects) and the γ's (contextual effects). The δ's are, in this analysis, ancillary parameters of interest. Correlated effects may emerge if ε_{it} contains firm- or network-specific effects. We discuss below how we deal with network or firm fixed effects, if present.

While using consumption data (a household, rather than an individual variable) creates additional complications, it also makes identification possible using network structure. This is because husbands and wives who work in different firms have their own distinct network of co-workers. This means that instead of dealing with a series of isolated networks (firms), we can generate links (or "edges/bridges") across networks precisely through spouses working at different firms. In other words, if our definition of peer was a co-worker and we were dealing with single households, identification would be impossible to achieve. On the other hand, if researchers were interested in studying spouse-specific outcomes (such as wages or hours), it would be possible to use within-household differences to identify network effects under the assumption that correlated effects (including time-varying ones) are household specific.

4.1. *Technical discussion*

4.1.1. An introductory example. As extensively discussed in the literature (De Paula, 2017), model (8) is not identified. To see our identification strategy, consider Figure 2.

Couples (A, B, D, E, J) are in circles, while co-workers are in the dashed rectangles (representing firms). For example, the husband in couple A and the husband in couple B work in firm f_1; the wife in couple J and the husband in couple D work in firm f_4, etc. The equivalent of equation (8) for household A in this simple network is:

$$\begin{aligned}\Delta \ln C_A = \delta_0 + \theta_w \Delta \ln C_D + \theta_h \Delta \ln C_B + \delta_w \Delta X_A^w + \delta_h \Delta X_A^h \\ + \gamma_w \Delta X_D^w + \gamma_h \Delta X_B^h + \varepsilon_A. \qquad (9)\end{aligned}$$

To make the sharpest possible case, assume that the only exogenous characteristics affecting consumption growth are some firm-specific shocks T_{f_j}. Hence, $\Delta X_A^h = T_{f_1}$, and so on. Previous evidence (Guiso *et al.*, 2004; Fagereng *et al.*, 2016) shows that firms pass onto wages some of

11. We ignore the complications related to non-unitary household consumption behaviour, although we acknowledge that in principle differences between θ_w and θ_h (or γ_w and γ_h) could reflect the different bargaining weights of the spouses in the intra-family consumption allocations.

their permanent value added shocks. For this reason, it is likely that some of the uninsurable shocks that shift household consumption originate from such firm-related shocks.[12]

Assuming the firm shock impacts the consumption of workers within the firm equally (just for simplicity), equation (9) rewrites as:

$$\Delta \ln C_A = \delta_0 + \theta_w \Delta \ln C_D + \theta_h \Delta \ln C_B + (\delta_h + \gamma_h) T_{f_1} + (\delta_w + \gamma_w) T_{f_2} + \varepsilon_A. \tag{10}$$

Since the corresponding consumption growth equations for peers D and B can be written as:

$$\Delta \ln C_D = \delta_0 + \theta_w \Delta \ln C_A + \theta_h \Delta \ln C_J + (\delta_h + \gamma_h) T_{f_4} + (\delta_w + \gamma_w) T_{f_2} + \varepsilon_D$$

and

$$\Delta \ln C_B = \delta_0 + \theta_w \Delta \ln C_E + \theta_h \Delta \ln C_A + (\delta_h + \gamma_h) T_{f_1} + (\delta_w + \gamma_w) T_{f_3} + \varepsilon_B$$

it is easy to see that one can use T_{f_3} and T_{f_4} as possible valid instruments for $\ln C_D$ and $\ln C_B$ in (10). A shock faced by firm f_3 affects the wage of the wife in couple B, and hence the consumption of this couple. This changes the consumption of couple A through the network effect (if $\theta_h \neq 0$), but (importantly) not through the common firm effect shared by A and B (firm f_1). Similarly, a shock faced by firm f_4 changes the wage of the husband in couple D, and hence their consumption. This changes the consumption of couple A through the network effect (if $\theta_w \neq 0$), but not through the common firm effect shared by A and D (firm f_2). Note also that, as long the as firm shocks exhibit some idiosyncratic variation, the instruments remain valid in cases in which all firms in the network are subject to some common shocks. These common shocks are, effectively, controlled by having the *own* firm shocks included in equation (10).

In the more general case in which consumption growth depends on firm-specific shocks and demographics, two sets of instruments become available: distance-3 peers' firm-related variables and average demographics of distance-3 peers. The use of an exogenous source of variation like firm effects makes our identification approach more transparent than the usual strategy based on the presence of an intransitive triad structure. In practice, it is similar to using experimental variation in distant nodes that percolate through the entire network (as long as networks effects are indeed present).

The example discussed in this section can be generalized to a case in which spouses' co-workers have separate endogenous and exogenous effects on household consumption growth. This case requires matrix notation but the intuition given in the example above carries through identically. To save space, we develop this case in the Appendix.

Equation (8) represents our main estimating equation. Note that first differencing log consumption has already eliminated certain types of household and individual fixed effects for the members of household i (for example, those arising from λ-constant effects, see Browning *et al.*, 1985). These fixed effects capture, among other things, unobserved heterogeneity in permanent income across individuals or households. They may also induce correlated effects in consumption levels if there is sorting on firms by workers of similar ability. For example, suppose that higher

12. It is easier to motivate this example by referring to "shocks" (*i.e.* unanticipated events). But as we make clearer below all we need is that certain variables (changes in firm-specific characteristics, anticipated or otherwise) are pre-determined as to satisfy an exclusion restriction. In other words, all we require is that—if consumption depends on firm-related variables—it depends only on those pertaining to the firms *we work for*, and not on those pertaining to the firms employing our distance-3 peers: the spouse of my co-worker (who is distance-3 from my spouse) and the spouse of the co-worker of my spouse (who is distance-3 from myself).

permanent income workers sort into better firms (Abowd *et al.*, 1999). Since higher permanent income workers also consume more, it is not surprising that their consumptions may be correlated even in the absence of any social influence. First differencing eliminates this type of correlated effects.[13]

However, one may be worried that certain type of network effects may still be present despite first differencing. For example, all workers in a given plant may face a common shock due to poor firm performance that is not captured by the plant-level variables we already control for. Call $f_{i^h(t)}$ and $f_{i^w(t)}$ the plant-specific effects for husband and wife in period t, and assume the error term ε_{it} in (8) contains the difference in these plant-specific effects, $\left(\Delta f_{i^w(t)}+\Delta f_{i^h(t)}\right)$. Hence, it may still be possible to have unobserved network effects correlated with our instruments. We consider two approaches. In the first, we restrict our analysis to plant stayers (*i.e.* estimation of (8) excludes household i if there is at least one member who is a plant mover). If the network effect is constant over time ($f_{i^j(t)}=f_{i^j}$ for $j=\{h,w\}$ and all t), first differencing eliminates the plant-related effects for those who do not change employer.[14] Our second approach uses the whole sample and adds fixed effects for the "transitions" $\Delta f_{i^w(t)}$ and $\Delta f_{i^h(t)}$.[15] Other specifications that are robust to correlated effects are discussed in Section 5.2.2.

4.1.2. Remarks on data construction. In the empirical analysis, we focus on couples where both spouses work. This is not a strong restriction given the high female participation rate in Denmark (above 80% in our sample period). However, we do face a series of constraints when it comes to data construction. First, we need to exclude couples that work in the same plant. Second, when we deal with multi-worker firms (which is the norm), we have to choose whether to construct average peers' consumption using simple or weighted averages, where the weights might depend on occupation and education. Finally, we need to avoid "feedback network effects". Suppose that persons 1 and 3 work at firm j and their spouses 2 and 4 work at another firm k (this is not an unlikely case given the important role of job search networks, see Montgomery, 1991; Pistaferri, 1999). In our scheme, the consumption of 1+2 depends on the consumption of 3+4. The way we construct the instrument would imply using the exogenous characteristics, or firm-level shocks of 1+2 as instrument for the consumption of 3+4, which will violate the exclusion restriction condition. We make sure to discard those cases when constructing our instruments.

13. In principle, people may sort into firms also because of common preferences (*e.g.* impatience), which unlike permanent income induce unobserved heterogeneity in consumption *growth* rather than *levels*. Most of the Euler equation literature assumes that preference parameters do not vary across households, ruling out this form of sorting by assumption. Moreover, our regressions include an extremely rich set of demographic controls (including average characteristics of co-workers), which are likely to attenuate the bias induced by this form of sorting. Finally, most of the sociology and labour economics literature emphasizes ability or productivity rather than preferences when motivating sorting behaviour. Indeed, in his seminal paper Montgomery (1991) writes of the "inbreeding bias between workers of similar ability" when describing job search activities and outcomes. Nonetheless, if workers were to sort into firms on the basis of preferences, it would still be possible to obtain consistent estimates of network effects by taking double differences of household and peers' consumption growth rates. Since under the null of no preference sorting our first-difference estimator and the one based on double differences are both consistent but the latter is less efficient, it is possible to implement a simple Durbin–Hausman–Wu test for the null of no sorting on preferences. This test fails to reject the null at conventional levels (the p-value of the test is 12%).

14. Of course, mobility across plants may be endogenous, and for this reason one may need to control for selection into staying with the same employer. Unfortunately, we do not have powerful exclusion restrictions to perform this exercise.

15. Hence we assume stationarity, or $\Delta f_{i^j(t)}=\Delta f_{i^j(s)}$ for $j=\{h,w\}$ and all s,t.

4.1.3. Relationship to the literature. Our identification strategy combines two strands of the empirical literature on the identification of peer/network effects. The first uses random shocks to a part of the network; the second uses fixed effects and rich controls to deal with the endogeneity concerns. Similarly to Kuhn *et al.* (2011), we employ a partial population experiment (Moffitt, 2001). However, while Kuhn *et al.* (2011) use a random lottery win to neighbours, we use random firm-level shocks to distance-3 peers. Bertrand *et al.* (2000) employ a rich set of controls and fixed effects to estimate the network effect of welfare take-up. In particular, the authors use local area and language group fixed effects, as well as a rich set of individual controls. In a similar spirit, we control in our analysis for individual, area, and year fixed effects, adding then a rich set of controls at the individual level.

5. RESULTS

5.1. *Network statistics*

Before presenting the estimation results, we provide some descriptive statistics on the network data. It is useful to recall the structure of the network we create. We start by selecting households where both husband and wife work ("household network line"). Their distance-1 peers are their co-workers ("firm network line"). Their distance-2 peers are the spouses of their co-workers (distance 1), if they are married and if the spouses work ("household network line" again). Their distance-3 peers are the co-workers of the spouses of their co-workers ("firm network line" again). Note that when we move along the household network line we are bound to get fewer nodes than when we move along the firm network line, simply because people can only have one spouse, but they can have multiple co-workers.

We consider several definitions of a co-worker. Our baseline definition takes individuals working in the same plant and weights more those with a similar occupation *and* level of education. The weights are constructed as follows. We first allocate individuals to five education groups (compulsory schooling, high school degree, vocational training, college degree, and Master's degree or PhD), $E=\{1\dots5\}$, and three occupation groups (blue collar, white collar, and manager), $O=\{1,2,3\}$. Next, we define a variable called "degree of separations" between any two individuals i^s and j^m as $d_{i^s j^m}=\left(\left|E_{i^s}-E_{j^m}\right|+\left|O_{i^s}-O_{j^m}\right|\right)$. Hence if i^s is a blue collar high school dropout ($E_{i^s}=2, O_{i^s}=1$) and j a college graduate manager ($E_{j^m}=5, O_{j^m}=3$), $d_{i^s j^m}=5$. For individuals with the same education and occupation, $d_{i^s j^m}=0$. We then create a quadratic weight variable

$$\omega_{i^s j^m}=\left(d_{i^s j^m}+1\right)^{-2} \tag{11}$$

and use it to generate weighted sums and averages. For example, household i wife's average consumption peers is given by:

$$\overline{\ln C_{it}^{w}}=\left(\sum_{j^m, j^m\neq i^w}\omega_{i^w j^m}\right)^{-1}\sum_{j^m, j^m\neq i^w}\omega_{i^w j^m}\ln C_{jt},$$

where j^m is the member m of family j ($m=\{h,w\}$) working in the same firm as i^w. We adopt a similar weighting procedure for the creation of the contextual variables.

Using a weighted adjacency matrix serves two purposes: (a) some nodes might be more "influential"in affecting behaviour; (b) they add variation to our right hand side variable. The use of a similarity index is also consistent with the homophyly literature (Currarini *et al*., 2011). Since weight assignment is inherently arbitrary, in the Robustness section we present results under alternative weighting procedures.

TABLE 3
Weighted network statistics (quadratic weights)

	Co-workers' group			**Distance-3 peers' averages**	
	Peers	Std. dev.	Variable	Wife's peers	Husband's peers
Distance 1			Age	38.62	38.46
Husband	73.30	179.77		(2.48)	(2.51)
Wife	95.07	233.34	Years of schooling	11.89	11.67
				(1.26)	(1.29)
Distance 2			Share of females	0.45	0.62
Husband	89.27	212.77		(0.21)	(0.20)
Wife	118.78	285.88	Share of blue collars	0.42	0.40
				(0.27)	(0.27)
Distance 3			Share of white collars	0.30	0.35
Husband	12,535	33,284		(0.19)	(0.21)
Wife	14,871	40,535	Share of managers	0.29	0.25
				(0.23)	(0.22)
	Co-workers' averages		# Kids 0–6	0.28	0.28
				(0.09)	(0.09)
Variable	Wife	Husband	# Kids 7–18	0.46	0.47
Age	38.24	38.59		(0.13)	(0.14)
	(5.61)	(5.62)			
Years of schooling	11.69	11.81	Firm size (in 1,000)	1.45	1.56
	(1.86)	(1.82)		(1.13)	(1.39)
Share of females	0.71	0.27	Firm size growth	0.001	0.001
	(0.26)	(0.26)		(0.002)	(0.002)
Share of blue collars	0.36	0.47	Public sector	0.53	0.63
	(0.37)	(0.41)		(0.30)	(0.30)
Share of white collars	0.39	0.22	Publicly traded	0.36	0.29
	(0.35)	(0.28)		(0.28)	(0.27)
Share of managers	0.25	0.31	Limited liability	0.02	0.01
	(0.34)	(0.36)		(0.08)	(0.08)
# Kids 0–6	0.28	0.27	Other	0.62	0.70
	(0.21)	(0.21)		(0.29)	(0.27)
# Kids 7–18	0.50	0.47			
	(0.30)	(0.29)			

Our networks span the entire Danish economy (or, more precisely, the part of the Danish economy that is observed working in firms). Looking at Table 3, we note that husbands have on average about 73 distance-1 (weighted) peers (or co-workers), while wives tend to work in larger firms (or in the public sector), with an average distance-1 peer network size of 95 (weighted) co-workers. When looking at distance 2 peers, the numbers are only slightly larger (average sizes are 90 and 119 for husbands and wives, respectively). Finally, to find distance-3 peers we again move along the firm network line and reapply the appropriate weighting scheme. Since wives have on average 119 co-workers and 95 of them have valid nodes (spouses who work), the expected number of distance-3 peers is therefore around 11,500. In practice, there are slightly more (around 14,870) due to a long right tail effect induced by skewness in firm size. In principle, the farther we move from the centre, the larger the network size. In practice, this is bounded by the size of the economy.

In the remaining part of Table 3, we present the average characteristics of co-workers and of distance 3-peers. As one would expect these characteristics are in line with the characteristics of the population in Table 1, as there is nothing specific about being a distance-2 or -3 node.

Identification of the parameters of interest relies upon variation in two main blocks: (i) changes in the composition of the workforce identified as distance-3 peers, in terms of their average age, education, gender, and so on; and (ii) economic shocks to distance 3 peers' workplace, *i.e.*, growth in the number of employees, changes in the growth rate (to focus on shocks, instead of random growth rates), and a change in the firm type (such as transition from a private equity to a publicly traded company or a process of privatization of a government-owned firm). Since we run first difference regressions and control for time, industry, and region effects in our first stage, these variables can be interpreted as "idiosyncratic shocks" to the firm's characteristics (in particular, its growth rate).

It is worth pointing out that while individual characteristics are quite similar in the general population and at the various distance nodes (indicating lack of significant sorting of "different" people in larger firms), for the way we identify distance-3 nodes, there is a higher likelihood of observing them in larger firms, potentially including the public sector (which employs during our sample period almost 50% of our population, as can be seen from the bottom right part of Table 3). We can also notice that our firm-level IV's identify firms that have faster growth than the whole population, are more likely to be located in the public sector, and are less likely to be limited liability companies.

5.2. *Euler equation estimates*

The main specification we adopt follows from the Euler equation (8):

$$\Delta \ln C_{it} = \delta_0 + \theta_w \Delta \overline{\ln C^w_{it}} + \theta_h \Delta \overline{\ln C^h_{it}} + \gamma_w \Delta \overline{X^w_{it}} + \gamma_h \Delta \overline{X^w_{it}} + \delta_w \Delta X^w_{it} + \delta_h \Delta X^h_{it} + \varepsilon_{it},$$

where C is household real consumption per-adult equivalent. We use the Luxembourg Income Study (LIS) equivalence scale, *i.e.*, $\sqrt{n_{it}}$ where n_{it} is family size. The set of exogenous characteristics (X) include: household controls (dummies for region of residence, number of children aged 0–6, and number of children aged 7–18), individual controls (age, age squared, years of schooling, dummies for blue collar, white collar, manager, industry dummies, a public sector dummy, firm size, firm-specific change in employment, and firm legal type), separately for husband and wife. We use the following contextual controls: age, age squared, years of schooling, number of children aged 0–6, number of children aged 7–18, share of female peers, share of blue collars, white collars, managers. All specifications also include year dummies.[16] We consider two sets of instruments. The first set consists of weighted averages of demographic characteristics of distance-3 peers: age, age squared, years of schooling, share of women, share of blue collars, share of white collars, share of managers, kids aged 0–6, and kids aged 7–13 ("Demographic IV's"). The second set of instruments includes firm-specific variables of distance-3 peers: firm size, firm-specific change in employment, firm type dummies, and a dummy for whether the firm is part of the public sector ("Firm IV's"). All these variables are then expressed in first differences.

The first three columns of Table 4 present estimates from three different specifications in which we control separately for the husband's and wife's networks. The other three columns re-estimate the same specifications imposing the assumption of a joint household network. Throughout the analysis standard errors are double clustered, with clusters defined by plant/occupation/education for both husband and wife.

In column (1), we present a standard OLS analysis on the first differenced consumption data. There are significant consumption network effects, which are similar for both husband and wife. However, these estimates are subject to the usual endogeneity (and reflection) problems. Indeed, Durbin–Hausman–Wu exogeneity tests (reported in column (2) and (3) for different combination of instruments) reject the null of exogeneity at conventional levels. In column (2), we thus present IV regression estimates using both demographics and firm-specific instruments, while in column

16. See the Appendix for more details on variable definitions.

TABLE 4
Baseline results

Model	(1) OLS FD	(2) IV FD	(3) IV FD	(4) OLS FD	(5) IV FD	(6) IV FD
Wife's peers ln C	0.11***	0.26**	0.30*	–	–	–
	(0.002)	(0.113)	(0.163)	–	–	–
Husband's peers ln C	0.13***	0.38***	0.37**	–	–	–
	(0.003)	(0.111)	(0.182)	–	–	–
Avg. peer's ln C	–	–	–	0.12***	0.32***	0.33***
	–	–	–	(0.002)	(0.061)	(0.078)
Demographic IV's	–	YES	NO	–	YES	NO
Firm IV's	–	YES	YES	–	YES	YES
p-value test $\theta_w = \theta_h$	0.000	0.510	0.819	–	–	–
p-value exogeneity test	–	0.002	0.015	–	0.001	0.004
p-value OID test	–	0.001	0.216	–	0.001	0.280
F-stat first stages						
Wife	–	44.14	78.27	–	–	–
Husband	–	43.14	59.85	–	–	–
All	–	–	–	–	62.84	111.40
Number of obs.	2,671,889	2,671,889	2,671,889	2,671,889	2,671,889	2,671,889

Notes: *,**,***= significant at 10%, 5%, 1%. Standard errors are double clustered for husbands and wives at the workplace, occupation, and education level. Dependent variable: Log of adult equivalent consumption. Individual controls (separately for husband and wife): Age, Age sq., Years of schooling, Occupation dummies, Industry dummies, Public sector dummy, Firm size, Firm growth, Firm type dummies. Household controls: Region dummies, # kids 0–6, # kids 7–18. Contextual controls (peer variables for husband and wife): Age, Age sq., Years of schooling, # kids0–6, # kids 7–18, share of female peers, shares of peers by occupation. Demographic IV's: Age, Age sq., Years of schooling, # kids0-6, # kids 7-18, share of female peers, shares of peers by occupation. Firm IV's: Public sector dummy, Firm size, Firm growth, Firm type dummy. We also control for year fixed effects. For details on the weighting schemes see the main text.

(3) we use only firm-level instruments. We also present first-stage statistics, which are generally much larger than conventionally acceptable thresholds (even discounting for the unusually large sample sizes).[17] Finally, since we have more instruments than endogenous variables, we also report p-values for overidentifying restriction tests (OID).

Our preferred specification is the one in column (3), where we only rely on more economically meaningful firm-level instruments and the OID test reveals no sign of misspecification. In the IV case, we continue to find non-negligible consumption network effects. In column (3), the husband's network effect is 0.37 and statistically significant at the 5% level, while the wife's network effect is slightly smaller, 0.3 and significant at the 10% level. It is important to quantify these effects. A 10% increase in the average consumption of the wife's (husband's) peers would increase household consumption by 3% (3.7%). However, given network size, this is a fairly aggregate shock—it is equivalent to a 10% simultaneous increase in the consumption of *all* peers (95 and 73 weighted peers, respectively for wife and husband—see Table 3). A different (and perhaps more meaningful) way of assessing these effects economically is to ask by how much household consumption would increase in response to a 10% increase in the consumption of a *random* peer in his/her network. We estimate this to be 0.03% in the wife's case and 0.05% in the case of the husband's.[18]

17. Note that, as remarked by Sanderson and Windmeijer (2016), in a setting with multiple endogenous variables significant first-stage F-statistics are necessary, but not sufficient, for identification of the parameters of interest. They propose the use of a conditional first-stage F-test statistic. Since we fail to reject the null hypothesis of a single endogenous variable, we do not consider this adjustment here.

18. This calculation is obtained by multiplying the estimated network effect by the probability of a member of the network experiencing the consumption increase (*i.e.* for the wife the effect is computed as $0.0032 = \frac{0.3}{95.1}$, while for the husband is $0.0050 = \frac{0.37}{73.3}$).

In monetary terms and evaluated at the average level of consumption, a 10% increase in the consumption of a random peer on the husband's side (corresponding to about \$5,000) would increase household consumption by about \$25 (and \$1,825 in the aggregate). On the wife's side, the effect would be \$15 (and \$1,425 in the aggregate). Since individual and aggregate effects may be very different, in Section 6 we attempt to quantify the macroeconomic implications of the network effects we estimate by simulating a number of policy scenarios.

Two things are worth noting. First, in columns (1)–(3), the effects of the husband's and wife's network on household consumption are economically very similar. In fact, when we test for the equality of the coefficients between husband and wife ($H_0 : \theta_h = \theta_w$) we cannot reject the null hypothesis of equality (with large *p*-values) in the more relevant IV models. Given this evidence, in the rest of Table 4 we re-estimate all models imposing that husband and wife belong to a single joint network (columns (4)–(6)), while acknowledging that our test may be biased towards the null due to the high standard errors induced by the instruments' low power. On the other hand, in our preferred specification of column (3) the estimates for husband and wife effects are economically very similar. The second aspect of the analysis that is worth highlighting is that OLS estimates are smaller than IV estimates. This is surprising given that a pure endogeneity bias would bias OLS estimates upwards. However, measurement error in peers' consumption may induce a bias that goes in the opposite direction, and it may be larger (in absolute value) than the endogeneity bias. Recall that our OLS specification has already eliminated a lot of the bias coming from observed and unobservable heterogeneity by first differencing and by the inclusion of a large set of socio-demographic controls. Hence, OLS estimates are more likely to reflect the downward bias of measurement error than the endogeneity upward bias. As an informal check, we re-estimated the OLS model without controlling for covariates and find larger estimates of $\widehat{\theta_h} = 0.5$ and $\widehat{\theta_w} = 0.45$. Moreover, any measurement error in levels is exacerbated by first differencing the data. Furthermore, it is possible that OLS is downward bias due to the exclusion bias highlighted by Caeyers and Fafchamps (2016). This is similar to the downward bias of auto-regressive coefficients in a time-series context (which can be interpreted as a circular network).[19]

Imposing a joint household network (as in columns (4)–(6)) results in predictably more precise IV estimates. The joint network effect is estimated to be around 0.33 and is statistically significant at the 1% level. The first-stage F-statistic in our preferred specification (column (6)) is 111, which shows that our instruments have strong identifying power. Diagnostic tests are similar to the ones obtained when imposing different networks. The economic interpretation of the joint network effect is also similar to the one presented above. A random peer's 10% increase in consumption would increase household consumption by 0.04%.

5.2.1. Heterogeneity of network effects. Are network effects heterogeneous? For example, one may believe that effects vary with network size: peer effects may be much more important in a small firm than in a large firm where personal and social contacts can be more diluted. Moreover, peer effects may depend on observables such as (weighted) network size,

19. Our IV procedure is also robust to the possibility of type I error (excluding workers who are peers) and type II error (including workers who are not peers) in the definition of the peer group. To see why, take Figure 2. We assume that the consumption of household A depends on the consumption of peer households B and D, and hence use as instruments the shocks of the firms employing the spouses of A^h and A^w's co-workers, *i.e.*, f_3 and f_4. But suppose that the "true"peers of household A are households B and K (a household totally outside the picture, whose members work at firms f_5 and f_6, say). Our instrument idea remains valid: we use f_3 and f_4 as instruments, but should rather use f_3, f_5 and f_6. As long as some co-workers are at least genuine peers, we may have a problem of low power of instruments (which is testable), but not a problem of violation of the exclusion restriction.

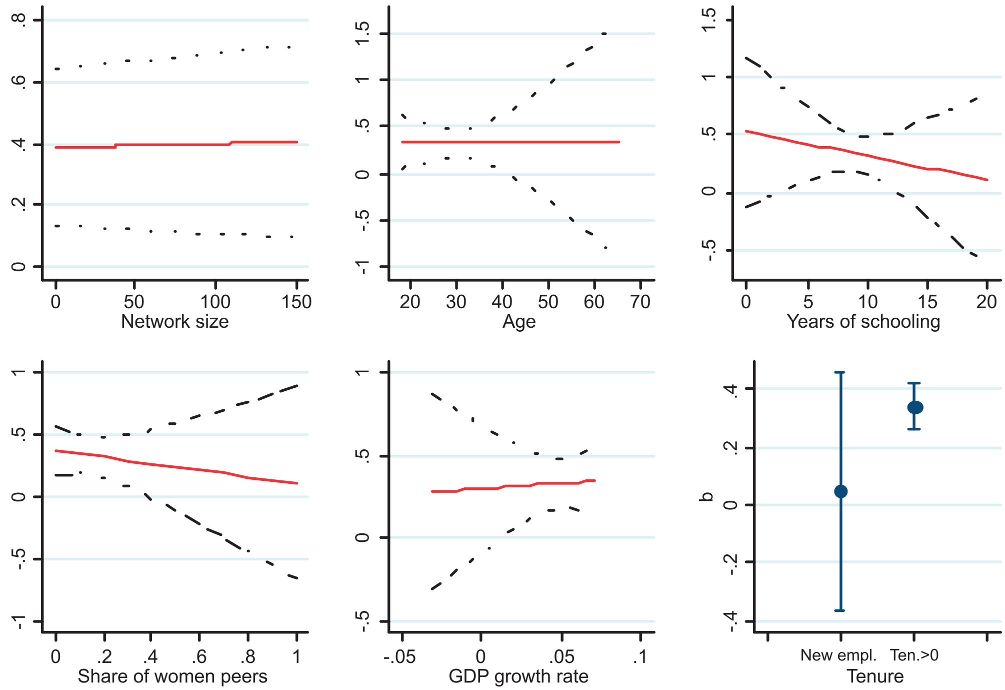

FIGURE 3

Heterogeneous network effects.

education, share of women, the business cycle, etc.. The effect of tenure is particularly important, as one may test whether social pressure increases with the time spent with a co-worker (effects may be small at low levels of tenure and larger at high levels of tenure).

Figure 3 shows how the consumption network effect varies with observable characteristics (network size, age, years of schooling, share of women peers, a measure of the business cycle, and tenure). These effects come from IV regressions similar to those in column (6) of Table 4, but with the inclusion of an interaction of the peers' average log consumption with the relevant source of heterogeneity. The graph also plots the upper and lower bounds of a 95% confidence interval. All demographic characteristics refer to the husband (due to his primary earner role). The regression estimates used in Figure 3 are reported in full in Table A2 in the online Appendix. Most of the effects we report are measured imprecisely, due to the fact that power of instruments dilutes when slicing the sample through interactive effects.

The first interesting result is that consumption peer effects do not vary significantly with the network size. Consumption network effects seems also stronger for the low-educated and in male-dominated professions. Interestingly, network effects are larger when the economy is booming and smaller during recessions—but estimates are significant only for a growth rate above 2% or so. Finally, we look at the effect of tenure. Unfortunately, our measure of tenure is limited and subject to left-censoring. We hence use a simple dummy for new employee. As expected network effects are present among those who have been in the firm for at least some time, but are statistically and economically absent for newly hired employee, perhaps because these individuals have yet to "learn" about co-workers' consumption choices and habits.

TABLE 5
Robustness to correlated effects

Model	(1) Baseline	(2) Stayers	(3) Transition FE	(4) Local shocks	(5) Sector shocks
Avg. peer's ln C	0.33***	0.31***	0.29***	0.44***	0.33***
	(0.078)	(0.106)	(0.063)	(0.049)	(0.073)
F-stat first stage	111.40	59.09	58.91	220.9	127.7
Number of obs.	2,671,889	1,323,086	2,671,889	2,671,889	2,671,889

Notes: *,**,***= significant at 10%, 5%, 1%. Standard errors are double clustered for husbands and wives at the workplace, occupation, and education level. Dependent variable: Log of adult equivalent consumption. Individual controls (separately for husband and wife): Age, Age sq., Years of schooling, Occupation dummies, Industry dummies, Public sector dummy, Firm size, Firm growth, Firm type dummies. Household controls: Region dummies, # kids 0–6, and # kids 7–18. Contextual controls (peer variables for husband and wife): Age, Age sq., Years of schooling, # kids 0–6, # kids 7–18, share of female peers, and shares of peers by occupation. Firm IV's: Public sector dummy, Firm size, Firm growth, and Firm type dummy. We also control for year fixed effects. The regression in column (1) is the baseline (from Table 4, column 6). For details on the other specifications, see the main text.

5.2.2. Other concerns. In this section, we investigate a variety of concerns and present some robustness analyses.

Correlated effects. A first concern with our estimates is the possibility that the error term may still contain network (correlated) effects which may generate spurious evidence of endogenous effects. To address this issue, we follow two strategies. In Table 5, column (2) we focus on a sample of firm stayers, for whom firm fixed effects are differenced out so that they are no longer a concern (column (1) reproduces, for comparison, the results of our preferred specification, that of Table 4, column 6). In column (3), we use the whole sample but include fixed effects for all possible cross-firm transitions (and assume stationarity).[20] Reassuringly, looking at stayers or including transition fixed effects leaves the results very similar to the baseline.

Bias from correlated effects may also come from co-workers suffering similar aggregate shocks (not already controlled for by the rich set of covariates we include in the baseline specification). Columns (4) and (5) are designed to further address these concerns. In column (4), we control for neighbourhood-specific shocks (measured by a full set of region-year dummies), while in column (5) we control for sector-specific shocks (measured by a full set of sector-year dummies). The results remain qualitatively similar.

Measurement error. Our measure of consumption, based on a budget accounting, may miss capital gains and capital losses, *i.e.*, it may fail to be accurate at the top and bottom of the consumption distribution. In column (2) of Table 6, we present results obtained using a measure of consumption that drops the top and bottom 1% of the consumption distribution. The estimate is essentially unchanged both in magnitude and statistical precision. Another way to address this issue is to focus on samples for whom capital gains and losses are unlikely to be important. Hence in column (3) we exclude stockholders, while in column (4) we focus on renters. In both cases, the estimate remains in the ballpark of the estimate obtained in the whole sample, which is comforting.

Weighting scheme. As emphasized above, the way we weight co-workers within the firm to form peer groups is inherently arbitrary. In columns (5)–(7) of Table 6, we assess whether the

20. In other words, we include dummies for each possible transition between any two firms in our data. While we do not impose the restriction that transitions from firm *A* to *B* are the same as transitions from *B* to *A*, we have to impose that the transition effects are constant over time (stationarity).

TABLE 6
Robustness to measurement error and weighting scheme

Model	(1) Baseline	(2) 1% trim	(3) Drop stockh.	(4) Renters	(5) Linear w.	(6) Exp. quadr.	(7) Sharp	(8) Placebo
Avg. peer's ln C	0.33***	0.32***	0.36***	0.24***	0.62***	0.30***	0.44***	0.01
	(0.078)	(0.066)	(0.088)	(0.082)	(0.164)	(0.082)	(0.083)	(0.028)
F-stat first stage	111.40	101.90	82.38	34.69	127.03	77.67	88.57	–
Number of obs.	2,671,889	2,588,299	2,016,137	318,317	2,671,889	2,671,823	2,171,426	2,671,889

Notes: *,**,***= significant at 10%, 5%, 1%. Standard errors are double clustered for husbands and wives at the workplace, occupation, and education level. Dependent variable: Log of adult equivalent consumption. Individual controls (separately for husband and wife): Age, Age sq., Years of schooling, Occupation dummies, Industry dummies, Public sector dummy, Firm size, Firm growth, and Firm type dummies. Household controls: Region dummies, # kids 0–6, and # kids 7–18. Contextual controls (peer variables for husband and wife): Age, Age sq., Years of schooling, # kids 0–6, # kids 7–18, share of female peers, and shares of peers by occupation. Firm IV's: Public sector dummy, Firm size, Firm growth, and Firm type dummy. We also control for year fixed effects. The regression in column (1) is the baseline (from Table 4, column 6). For details on the other specifications, see the main text.

results are robust to adopting different weighting schemes (and report the main features of the networks in Tables A3–A5 in the online Appendix). In column (5) we experiment with a linear scheme, while in column (6) we consider a richer quadratic weighting scheme based on education, occupation, and age (to capture tenure effects). To keep the number of groups within the feasible range, we consider just two age ranges, 40 and less, and more than 40. Finally, in column (7) we use a sharper weighting scheme in which all workers in the same plant and occupation are treated equally (regardless of their education). The estimates of the endogenous effects vary in size across adopted schemes. However, once these estimates are appropriately rescaled for the larger or smaller peer group, the elasticities we obtain are in the same ballpark as those in column (1), and discussed above. For example, a 10% increase in the consumption of a random peer produces a 0.05% effect for the linear scheme in column (2), a 0.053% for the scheme of column (3), and finally a 0.036% effect for the same occupation scheme in column (4) (as opposed to 0.04% in the baseline specification).

In column (8) we consider a "placebo" weighting scheme. This exercise is motivated by the consideration that our results could still be spurious if there are some unobserved factors running through the economy which produce correlation in consumption patterns that have nothing to do with network effects (and that may originate from non-linear aggregate effects not perfectly controlled by year dummies). To assuage these fears, we construct placebo samples where we randomly assign workers to firms, keeping firm sizes constant. In other words, we scramble the firm identifier but keep the number of workers at each firm identical as in the actual dataset. By construction, this eliminates any form of sorting. We then recompute peer consumption as explained in Section 5.1. The results, based on 50 replications, are reported in the last column of Table 6. They show that the main estimated effects are not spurious. When individuals are randomly allocated peers, their consumption is independent of that of their randomly allocated peers, with an estimated small network effect of 0.01 and a large standard error.

Responding to labour supply or saving changes?. Our definition of consumption from the administrative data is the difference between income and saving. Since $\ln C = \ln((Y-T)-\Delta A) \cong \ln(Y-T)-s$, (where $s=\Delta A/(Y-T)$ is the household's saving rate), one can decompose the consumption response to peers' consumption into two parts: a response to a change in peers' disposable income (or, broadly, labour supply) (θ^Y), and a response to a change in peers'

TABLE 7

Response to effort versus saving rate

Avg. peer's ln Y	0.09
	(0.166)
Avg. peer's saving rate	0.33***
	(0.078)
Number of obs.	2,671,889

Notes: *,**,***= significant at 10%, 5%, 1%. Standard errors are double clustered for husbands and wives at the workplace, occupation, and education level. Dependent variable: Log of adult equivalent consumption. Individual controls (separately for husband and wife): Age, Age sq., Years of schooling, Occupation dummies, Industry dummies, Public sector dummy, Firm size, Firm growth, and Firm type dummies. Household controls: Region dummies, # kids 0–6, and # kids 7–18. Contextual controls (peer variables for husband and wife): Age, Age sq., Years of schooling, # kids 0–6, # kids 7–18, share of female peers, and shares of peers by occupation. Firm IV's: Public sector dummy, Firm size, Firm growth, and Firm type dummy. We also control for year fixed effects. The regression in column (1) is the baseline (from Table 4, column 6). For details on the other specifications, see the main text.

saving rate (θ^s). This decomposition is informative about what kind of peer behaviour influences household consumption decisions most.

We implement this idea by running the following regression:

$$\Delta \ln C_{it} = \delta_0 + \theta^Y \Delta \overline{\ln Y_{it}} + \theta^s(-\Delta \overline{s_{it}}) + \gamma_w \Delta \overline{X_{it}^w} + \gamma_h \Delta \overline{X_{it}^w} + \delta_w \Delta X_{it}^w + \delta_h \Delta X_{it}^h + \varepsilon_{it}.$$

If peers' labour supply and saving rates decisions contribute in the same way to the overall response, we should find $\theta^Y = \theta^s$. In Table 7, we find $\theta^Y = 0.094$ (s.e. 0.17) and $\theta^S = 0.334$ (s.e. 0.08). It thus appears that both types of responses matter (although the income component is noisy), but that consumption is more likely to be affected by changes in peers' saving/borrowing than changes in peers' effort or labour supply. As we shall see, this result can be read as consistent with the main finding from the demand analysis below—namely, that consumers do not seem to over-react to "visible" or conspicuous consumption patterns of peers. In the context examined here, the response to labour supply/effort activities (which are presumably more visible than saving or borrowing decisions, at least in a workplace context) is actually smaller than the response to saving/borrowing choices of peers.[21]

Other robustness checks. We have performed additional robustness checks. For reason of space, they are discussed in the Appendix. None of these extra checks affect the qualitative pattern of results. Overall, we take the series of results presented in Tables 7 (and Table A2 in the online Appendix) as reassuring. The estimated effects appear robust to several potential sources of bias and change in predictable ways when we change the way we weight co-workers within a plant.

5.3. *Demand estimation*

The results presented in the previous section point to the presence of intertemporal distortions on consumer behaviour. Table 2 suggests that intertemporal distortions may also be compatible with

21. A different way to look at the importance of the labour supply effect is to regress the growth in disposable income against its peer equivalent (using the same instruments we use in the consumption case). If we do so, we obtain an estimate of 0.086 (s.e. 0.023). If labour supply was the the only channel at play, the association between individual and peer consumption should be weaker, not stronger (given that consumption is smoother than income).

the presence of intratemporal distortions, which may have very different policy implications, as well as suggesting different theoretical mechanisms.

In this section, we follow the structure developed in Section 3.1 and estimate demand equations for "visible" (V) and "non-visible" (I) goods. In particular, we run the following regressions:

$$\omega_{it}^{j} = X_{it}'\alpha_{0j} + \alpha_{1j}\overline{\ln C}_{it} + \beta_{0j}\ln C_{it} + \beta_{1j}\left(\ln C_{it} \times \overline{\ln C}_{it}\right) + \upsilon_{it}^{j} \tag{12}$$

for $j = \{V, I\}$. Neutral (N) goods (which we assume include only food at home) represent the excluded category. The detailed categorization of what we include in the three types of goods is presented in Table A6 in the online Appendix. As discussed in Section 3, we test whether the average consumption peer variables are insignificant determinants of the demand for goods (*i.e.* $\alpha_{1j} = \beta_{1j} = 0$ for all $j = \{V, I\}$), *controlling for* total spending $\ln C_{it}$. In practice, we faced some collinearity problems and hence we only report the results of a simpler specification in which we omit the interaction (and assume $\beta_{1j} = 0$ for all j). This allows us to perform a simple test of reshuffling, *i.e.*, testing that $\alpha_{1V}\alpha_{1I} < 0$.

We report results for two samples. The first sample is all households that can be matched with the tax registry (independently of marital and work status), comprised of 2,437 households (the "ALL" sample). We do not have distance-3 instruments for all households in this sample (as this depends on both the work and marital status) and hence we run simple OLS regressions. Our second sample is a perfect match with our tax registry baseline sample, and is hence much smaller given the restrictions we apply for estimation (454 households, or the "MATCH" sample). For these households, we can run IV regressions instrumenting peer consumption with distance-3 instruments as in the Euler equation case discussed above.

The results are reported in Table 8. In columns (1) and (2), we report estimates of (12) for the "ALL" sample. There is no evidence that conspicuous consumption changes intratemporal allocations. Controlling for total consumption, the marginal effect of peers' consumption, $\frac{\partial \omega_{it}}{\partial \overline{\ln C}_{it}}$, is small and statistically insignificant for both visible and non-visible goods. The estimates suggest that visible goods are luxuries and non-visible goods are necessities. Note that the results do not depend on the richness of controls used, and are confirmed even when we have no controls in the regression besides total consumption and peer consumption. In columns (3) and (4), we replicate this regression on our "MATCH" sample, and are now able to instrument peer consumption with distance-3 exogenous firm-level shocks and characteristics. The results are qualitatively unchanged.[22] Interestingly, these findings are similar to those of Lewbel *et al.* (2018), who also conclude that peer effects are similar for visible and non-visible goods.[23]

There is some inherent arbitrariness in how we classify goods into visible, neutral, and non-visible categories. To counter this criticism, we disaggregate spending into all the 30 categories considered by Heffetz (2011), and run the budget share regression (12) separately for each good category (imposing again $\beta_{1j} = 0$ for all j). See Table A7 in the online Appendix for the exact category definition. Figure 4 plots the estimated coefficients (and corresponding 95% confidence intervals) against the degree of visibility as estimated in Heffetz (2011). We also plot the local regression line to detect any possible relationship between the visibility index and the estimated coefficients. We do this for the MATCH sample (where we can instrument peer consumption with distance 3 instruments). In principle, the regression coefficient should rise with the degree of

22. The results also remain qualitatively unchanged if we instrument log household consumption with log household income or with the registry-based consumption measure.

23. We also conduct the reshuffling test mentioned in footnote 8, but fail to reject the null of no reshuffling with a bootstrap p-value of 46% (based on 200 replications).

TABLE 8
Demand estimation

	ALL sample		MATCH sample	
	Visible cons. (1)	Not-visible cons. (2)	Visible cons. (5)	Not-visible cons. (6)
ln C	0.205***	−0.132***	0.270***	−0.165***
	(0.007)	(0.007)	(0.018)	(0.017)
Avg. peer's ln C	0.005	−0.000	0.004	0.007
	(0.011)	(0.011)	(0.028)	(0.027)
Observations	2,435	2,437	454	452

Note: *,**,***= significant at 10%, 5%, 1%. The dependent variables are budget shares for three consumption groups: Visible, Not-visible and Neutral. The omitted category is Neutral. Individual controls: Age, Age sq., Years of schooling, Occupation dummies, Industry dummies, Public sector dummy, Firm size, Firm growth, and Firm type dummies. Household controls: Year dummies, Region dummies, # kids 0–6, and # kids 7–18. Columns (1)–(4) are OLS estimates; in columns (5) and (6) the average peers' log consumption is instrumented using distance-3 instruments.

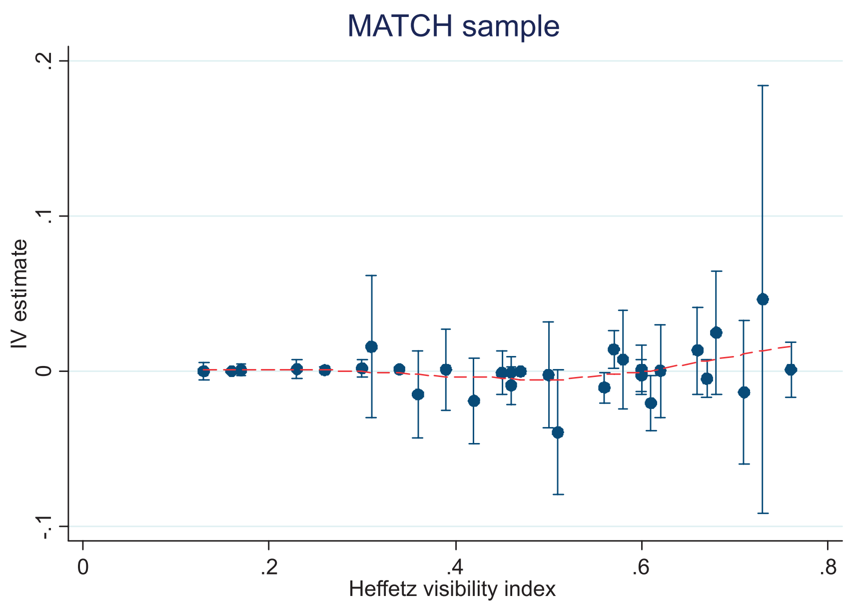

FIGURE 4
The relationship between the shift in budget shares due to peer consumption and the Heffetz visibility index.

visibility if there were any intratemporal effects. However, the disaggregated evidence is similar to the one noted above. The effect of peer consumption on the budget share on good j appears independent of the degree of conspicuousness of the good. The relationship is increasing only for highly conspicuous goods, but the estimates are very noisy.

5.4. *Risk sharing*

As discussed in Section 3.2, another reason for observing a correlation between individual and peer consumption is because of risk sharing within the firm. The theory of risk sharing states that when risks are shared optimally, consumption growth of two individuals who are part of a risk sharing agreement will move in lockstep even if the two individuals do not influence each other's

TABLE 9
Tests of risk sharing

	ALL sample	
	(1)	(2)
$\Delta \ln Y$	0.135*	0.131*
	(0.069)	(0.069)
$\overline{\Delta \ln Y}$		0.046
		(0.078)
Observations	2,454	2,454

Note: *,**,***= significant at 10%, 5%, 1%. Individual controls: Age, Age sq., Years of schooling, Occupation dummies, Industry dummies, Public sector dummy, Firm size, Firm growth, and Firm type dummies. Household controls: Year dummies, Region dummies, # kids 0–6, and # kids 7–18. The dependent variable is the difference between log consumption in the survey and log consumption in the tax records.

consumption. The extreme case is where co-workers only observe income but do not observe consumption (*i.e.* all relevant consumption is domestic). However, this is enough to generate risk sharing as long as we believe problems of private information or limited enforcement are more easily solved within the strict confines of the workplace.

Results reported in Table 4 already reject the strongest form of full insurance (*i.e.* that individual consumption should move at the same rate as aggregate peer consumption). The way we test for partial risk sharing is explained in Section 3.2. We consider the regressions:

$$\ln C_{it}^S - \ln C_{it}^T = X_{it}'\pi_0 + \pi_1 \Delta \ln Y_{it} + \pi_2 \Delta \overline{\ln Y_{it}} + \upsilon_{it}$$

and test whether $\pi_1 < 0$ and $\pi_2 > 0$.

The results are reported in Table 9. Since we are only interested in the sign of the relationship between the coefficient on income growth and the consumption differences between survey and tax registry data, here we focus only on the "ALL" sample and perform simple OLS regressions (results for the "MATCH" sample are similar qualitatively but noisier given the smaller sample size). Risk sharing would suggest a negative association between earnings growth and the survey-tax record consumption log-differential. This is because individuals who suffer a negative income loss should receive a transfer from peers, which would increase the survey-based measure of consumption (which includes the "shared" transfer) relative to the tax records measure (which does not). In the data there is actually a positive, weakly statistically significant association. Similarly, there is no evidence of a positive association between average earnings growth of peers and the survey-tax record consumption log differential. We conclude that it is unlikely that our results of significant peer effects are spuriously coming from risk sharing within the firm.

6. IMPLICATIONS: AGGREGATE EFFECTS

The effects of macroeconomic stabilization policies may depend on the presence of peer effects. Small stabilization policies may have larger or smaller effects than in a world where peer effects are absent because of social multiplier effects. Here, we discuss a simple macro experiment based on our empirical estimates. In this experiment, we neglect General Equilibrium effects on asset prices, labour supply, and so forth, to highlight the role of network effects in the sharpest possible way.

We start from the consideration that a tax/transfer imposed on a group may reverberate through the entire distribution, depending on the degree of connectedness of individuals. A "benchmark" multiplier, which abstracts from the degree of connectedness, is constant across individuals and equals 1.5 (obtained as $(\mathbf{I}-\widehat{\theta}\mathbf{ii}')^{-1}$, where $\widehat{\theta}\cong 0.33$ from the regression of Table 4, column 6), so aggregate effects may potentially be important. We should note that the specific multiplier is only valid in a world where the network is full, *i.e.*, all the nodes are directly connected, which is clearly not the case in a standard setting and in our specific application. We therefore have to account for the degree of connectedness as well as for the introductory point of the policy (*i.e.* which group is directly targeted) in order to understand the aggregate implications of network effects. To do so, we engineer a series of experiments, summarized in Table 10. We present the multiplier effects of a one-time policy (in a static framework) as well as several moments of the resulting distribution of consumption. We do this in order to understand the level and distributional effects of such policy experiments. In the first row of Table 10, we present the actual moments from the 1996 sample (our last sample year).

Our first three experiments consist of transferring the equivalent of 1% of aggregate consumption equally among: (a) households in the top 10% of the consumption distribution, (b) a 10% random sample of households, and (c) households in the bottom 10% of the consumption distribution. These three policies are financed by issuing debt and running a government deficit. As an alternative to a debt-financed policy, we consider: (d) a purely redistributive policy in which the receivers of the transfers are households in the bottom 10%, and the policy is financed by a "tax" to the top 10% of households. Alternatively, the government transfer is a consumption coupon, so it is entirely consumed. Note that we abstract from the possibility that marginal propensities to consume (MPCs) are heterogeneous (Jappelli *et al.*, 2014).

Consider the first experiment, which consists of distributing resources to households in the top 10% of the consumption distribution. In a world without network effects, this would increase aggregate consumption by 1%, with an implied multiplier of 1.01. With network effects, the implied multiplier effect is instead slightly larger, 1.012.[24] There is also a slight increase in the dispersion of consumption, as measured by the standard deviation of log consumption or the 90-10 percentile difference in log consumption. The reverberation effects are concentrated in the top half of the consumption distribution (as can be seen by looking at the 90/50 and 50/10 log consumption differences). What can be learned from this experiment is that consumption policies targeted at the top 10% of the consumption distribution (presumably also the wealthier households), have limited aggregate effects, and in particular do not spread along the distribution of consumption. The reason is that households at the top of the consumption distribution have fewer direct connections and their network structures are smaller and more sparse then those of random households in the population.

The next experiment (where we target a random 10% of households) confirms this intuition. In this case the multiplier is 1.017 and consumption inequality declines. A look at the 90-50 and 50-10 differences reveals that policies that target a random sample of households (most likely located in the middle of the distribution) have larger, and more far reaching consequences than policies of identical magnitude targeted at the top 10%. This is because those households have larger and denser networks than those at the top.

Even larger aggregate effects are found when the policy targets the bottom 10% of households. In this case the multiplier effect is noticeably larger than in the previous cases (1.034), with a much larger fall in dispersion (a 13% decline in the standard deviation of log consumption). These results suggest that households at the bottom of the distribution have larger and denser networks,

24. The multiplier is obtained as the ratio of the post-transfer to the pre-transfer aggregate consumption. We obtain post-transfer aggregate consumption using a simple iterative algorithm. After increasing the consumption of households targeted by the policy by the amount of the transfer, we compute peer consumption using the weighted formula explained in the text. We then compute the new level of consumption of all households in the sample using the network effect estimate. We recompute peer consumption, and so forth. We stop after 10 iterations.

TABLE 10
Counterfactual policy simulations

Transfer recipients	Implied multiplier	Std.Dev. log cons.	90-10 log difference	50-10 log difference	90-50 log difference
Baseline	–	0.729	1.5795	0.8630	0.7165
Top 10%	1.012	0.736	1.5818	0.8631	0.7187
Random 10%	1.017	0.718	1.5608	0.8506	0.7102
Bottom 10%	1.034	0.601	1.4525	0.7349	0.7176
Balanced budget	1.021	0.593	1.4216	0.7347	0.6869

which tend to be concentrated among households with similarly low consumption levels. Indeed, a look at the 50/10 and 90/50 percentile differences show that the latter barely moves (relatively to the baseline), while the former declines substantially.

In the final row of Table 10, we consider a balanced budget experiments in which a transfer to poorer households is financed by a tax imposed on the richer households (who hence mechanically reduce their consumption). This case yields an intermediate multiplier effect (1.021), with the largest reduction in dispersion among all experiments. This is because there are now richer effects: households connected with those at the top (presumably near the top themselves) reduce their consumption, while households connected with those at the bottom (presumably also located in the bottom half) increase it. The result is that the post-policy consumption distribution becomes more compressed, and the larger degree of connectedness at the bottom than at the top drives aggregate consumption upwards.

In conclusion, what we learn from the experiments detailed in Table 10 is that stimulus policies can have quite differential impacts on aggregate consumption and on the distribution of household consumption depending on the groups that are directly targeted (and their overall connectedness with the different segments of the population). Another important, and obviously related lesson is that the vector of social multipliers (generally computed as $(1-\theta)^{-1}\mathbf{i}$) can be misleading in the presence of fairly general network structures.[25] A related point based on social distance is given in Glaeser *et al.* (2003).

7. CONCLUSIONS

This article builds a consistent theoretical framework for consumption choices within and between periods that is able to capture social effects and allows us to distinguish between different ways in which social interactions can emerge. We take the testable empirical predictions that come from the model and bring them to bear on a very rich dataset containing (derived) information on the consumption of a large sample of the Danish population and the social networks they span (at the household and firm level). We find that peers' consumption affect intertemporal consumption decisions. We do not find evidence that network effects distort the demand for goods, although we have admittedly less power than when estimating Euler equations. As well, we find little evidence that peer effects emerge as a way of rationalizing risk sharing agreements. We have discussed the policy consequences of these results using simple stimulus policy experiments that transfer consumption resources to different groups in the population. The results highlight two important conclusions: the effects of the policies depend on the degree of connectedness of the group directly targeted, and the use of social multipliers can be misleading when network structures are not full.

Our results could be extended in a number of directions. While we have emphasized the importance of peers defined on the basis of co-worker relationships, it would be possible to

25. In our context, the vector of social multiplier effects is $(\mathbf{I}-\theta\mathbf{G})^{-1}\mathbf{i}$, where $\mathbf{G}$ is the (sparse) adjacency matrix.

construct family networks or location networks. Family identifiers, for example, could be used to match parents and children, or siblings. The problem with this approach is that a non-negligible number of households may be completely disconnected (*i.e.* older households or only children); moreover, there is the theoretically complex problems of separating altruistic behaviour from network effects. On the location side, we observe the municipality where the household resides. Here, we may face the opposite problems (*i.e.* the network may be too large and composed of people who do not interact socially in any meaningful way). Another interesting extension is to study whether peer effects have some dynamic effects that are different than inducing a response in the same year, as assumed here. On the theoretical side, it could be possible to test whether peers provide primarily "information". If the information story is an important one we should see it emerging mostly among goods with larger informational asymmetries (as reflected in pricing). Unfortunately, our microdata on spending are too limited to implement such exercise. Another fruitful area for future work is the importance of liquidity constraints or internal habits in shaping behaviour in the presence of network effects.

A. APPENDIX

A.1. *Additional robustness checks*

In Appendix online Table A2, we present the results of running additional robustness checks.

A concern we try to address in the main text is the possibility that the error term contains correlated effects generating spurious evidence of endogenous effects. Table 5, column (2), addressed this issue by focusing on a sample of firm stayers (*i.e.* the vector of observations on the dependent variable drops individuals who change employer; and we remove the rows of the matrix of covariates that correspond to those observations). The logic for this was explained in the concluding paragraph of Section 4.1.1. A different way of doing this exercise is to only use plants that experience no change in composition from 1 year to the next (*i.e.* no workers move out and no new workers join in). Since this is a fairly stringent criterion, it reduces the sample size to only 862,592 observations and result in a large decline in instrument power (an F-stat of 13). Nevertheless, the results remain qualitatively similar—see column (2) of online appendix Table A2. The point estimate is close to the baseline but due to the reduced power, it is less precise (a p-value of 13%).

In column (3), we add controls for age and education of the spouses in levels (rather than in first differences). The results are unchanged.

To complement the analysis discussed in Section 5.2.2, in column (4) we regress $\Delta\log(Y-T)$ on $\Delta\overline{\log(C)}$ using the same set of instruments as in the baseline specification. The estimate is small and significant -0.03 (s.e. 0.01). This is consistent with the evidence we reported in Table 7, since it shows that the main "channel" through which consumption network effects operate is saving, not income (labour supply). This is perhaps because of constraints to labour supply or other forms of inflexibility in work schedules, etc.

A.2. *Reshuffling*

Suppose that there are three types of goods, V ("visible"), I ("not visible"), and N ("neutral"). To see reshuffling with a simple example, assume a simplified functional forms for $a_t(.)$ and $b_t(.)$:

$$\ln a_t(.) = \alpha_0 + \sum_{k=\{V,I,N\}} \alpha_{0k} \ln p_t^k + \alpha_{1V}\overline{\ln C}_{it} \ln p_t^V + \frac{1}{2}\sum_{k=\{V,I,N\}}\sum_{j=\{V,I,N\}} \eta_{kj} \ln p_t^k \ln p_t^j$$

$$\ln b_t(.) = \sum_{k=\{V,I,N\}} \beta_{0k} \ln p_t^k + \beta_{1V}\overline{\ln C}_{it} \ln p_t^V$$

in which peers' consumption shifts only the visible consumption component of the price indexes. Moreover, assume for simplicity quasi-homotheticity ($\rho \to 1$). Then budget shares are:

$$\omega_{iVt} = \alpha_{0V} + \alpha_{1V}\overline{\ln C}_{it} + \sum_{k=\{V,I,N\}} \eta_{Vk} \ln p_t^k + \left(\beta_{0V} + \beta_{1V}\overline{\ln C}_{it}\right) \ln(C_{it}/a_t(.)) \tag{A.1}$$

$$\omega_{ijt} = \alpha_{0j} + \sum_{k=\{V,I,N\}} \eta_{jk} \ln p_t^k + \beta_{0j} \ln(C_{it}/a_t(.)) \tag{A.2}$$

for $j=\{I,N\}$.

To see why there is reshuffling, assume that peers' effects are positive ($\alpha_{1V}>0$). It is straightforward to show that $\frac{\partial \omega_{ijt}}{\partial \ln C_{it}}=-\beta_{0j}\alpha_{1V}\ln p_t^V$ for all $j=\{I,N\}$. If goods are normal, $\beta_{0j}>0$, and hence the demand for goods that are not visible declines as peers' consumption increases. But since budget shares sum to one (and hence $k=\{V,I,N\}\frac{\partial \omega_{ikt}}{\partial \ln C_{it}}=0$), the demand for the visible goods must increase. Hence, there is a form of "reshuffling" as peers' consumption increases: the demand for visible goods increases and that for goods that are not visible declines.[26]

A.3. *Discussion about identification*

Here, we generalized the example discussed in the main text. The multiple network case is also discussed elsewhere (*i.e.* Goldsmith and Imbens, 2013).

We allow the spouses' co-workers to have separate endogenous and exogenous effects on household consumption growth. This describes well our data, which are a combination of household level variables, *i.e.*, income and wealth (and therefore consumption), as well as individual level variables such as occupation, education, etc.

The model primitives are as follows:

- *Household Level Variables:* $\Delta\mathbf{c}$ is the $(N\times 1)$ vector of household changes in (log) consumption.
- *Individual Level Variables:*
    - $\Delta\mathbf{X}$ is a $(2N\times k)$ matrix of individual changes in exogenous characteristics. For simplicity of notation, we focus on the $k=1$ case. Just out of convention, we order the husband characteristics in each couple in the first N rows, followed by the wives' characteristics in each couple in the remaining N rows, *i.e.*, $\Delta\mathbf{X}=\begin{pmatrix}\Delta\mathbf{X}_h & \Delta\mathbf{X}_w\end{pmatrix}'$. However we note that the $\Delta\mathbf{X}$ can contain the firm-level exogenous shock.
    - Let also $\mathbf{S}_h$ ($\mathbf{S}_w$) be a transformation $(2N\times N)$ matrix that maps households into husbands (wives). Given our conventional ordering, $\mathbf{S}_h=\begin{pmatrix}\mathbf{I} & \mathbf{0}\end{pmatrix}'$ and $\mathbf{S}_w=\begin{pmatrix}\mathbf{0} & \mathbf{I}\end{pmatrix}'$. Hence $\mathbf{S}_h'\Delta\mathbf{X}$ (respectively, $\mathbf{S}_w'\Delta\mathbf{X}$) will be the vector of husband's (wife's) exogenous characteristics changes.
    - Let $\mathbf{D}$ be the $(2N\times 2N)$ social network at the person level. The generic element of $\mathbf{D}$ is:
    $$d_{i^l j^m}=\begin{cases}1 & \text{if } i^l \text{ connected to } j^m \text{ (for } l,m=\{h,w\}),\end{cases}$$
    where as before i^h and i^w denote husband and wife in household i, respectively, and $d_{i^l i^m}=0$ for $l,m=\{h,w\}$. The number of connections for a generic individual i^l is given by $n_{i^l}=\sum_{m=\{h,w\}}\sum_{j=1}^N d_{i^l j^m}$.[27]
    - Call $\mathbf{n}$ the $(2N\times 1)$ vector with generic element n_{i^l}. The row-normalized adjacency matrix is: $\mathbf{G}=\mathrm{diag}(\mathbf{n})^{-1}\mathbf{D}$ with generic element $g_{i^l j^m}=n_{i^l}^{-1}d_{i^l j^m}$.
    - Given this notation, $\mathbf{S}_h'\mathbf{G}(\mathbf{S}_h+\mathbf{S}_w)=\mathbf{G}_h$ is the husband-induced household network, with typical entry given by $\sum_{m=\{h,w\}}g_{i^h j^m}$, and identifies the households who are connected to the husbands (wives) of the N households in our sample (symmetrically, $\mathbf{S}_w'\mathbf{G}(\mathbf{S}_h+\mathbf{S}_w)=\mathbf{G}_w$ is the wife-induced household network). Hence $\mathbf{G}_h\Delta\mathbf{c}$ ($\mathbf{G}_w\Delta\mathbf{c}$) is the vector of husband's (wife's) peers' consumption growth.
    - Similarly $\mathbf{S}_h'\mathbf{G}\Delta\mathbf{X}$ ($\mathbf{S}_w'\mathbf{G}\Delta\mathbf{X}$) is the vector of the husband's (wife's) peers' exogenous characteristics changes.

Given this notation, the matrix equivalent of (8) can be written (omitting the constant terms for simplicity) as:

$$\Delta\mathbf{c}=(\theta_h\mathbf{G}_h+\theta_w\mathbf{G}_w)\Delta\mathbf{c}+(\mathbf{S}_h'\mathbf{G}\gamma_h+\mathbf{S}_w'\mathbf{G}\gamma_w+\mathbf{S}_h'\delta_h+\mathbf{S}_w'\delta_w)\Delta\mathbf{X}+\xi \tag{A.3}$$

26. It is possible that the price indexes depend on peers' *visible* (rather than *aggregate*) consumption. In this case, there will be a downward bias in the estimation of α_{1V} (a simple measurement error analogy). Unfortunately, we only observe average peers' total consumption, but not the average peers' consumption of visibles. One can show that (under the simplifying assumption $\beta_{1V}=0$, see De Giorgi *et al.*, 2016):

$$p\lim\widehat{\alpha}_{1V}=\alpha_{1V}B$$

$$p\lim\widehat{\alpha}_{1I}=\alpha_{1I}B,$$

where B depends on higher moments of the joint distribution of consumption and its components. Hence, it is still possible to construct a test of reshuffling based on $\widehat{\alpha}_{1V}\widehat{\alpha}_{1I}$, which converges to $\alpha_{1V}\alpha_{1I}B^2$. Under reshuffling, this product should be negative (as α_{1V} and α_{1I} move in opposite directions and $B^2\geq 0$).

27. A generalization of this is weighting the influence of different connections differently, *i.e.*, $\widetilde{n}_{i^l}=\sum_{m=\{h,w\}}\sum_{j=1}^N d_{i^l j^m}\omega_{i^l j^m}$. This is what we do in the empirical analysis.

If $(\mathbf{I}-(\theta_h\mathbf{G}_h+\theta_w\mathbf{G}_w))$ is invertible, we can use the Neumann series expansion of a matrix (Meyer, 2000, p. 527) to write:

$$(\mathbf{I}-(\theta_h\mathbf{G}_h+\theta_w\mathbf{G}_w))^{-1}=\sum_{k=0}^{\infty}(\theta_h\mathbf{G}_h+\theta_w\mathbf{G}_w)^k$$
$$=\mathbf{I}+(\theta_h\mathbf{G}_h+\theta_w\mathbf{G}_w)+(\theta_h\mathbf{G}_h+\theta_w\mathbf{G}_w)^2+\cdots \quad \text{(A.4)}$$

which is satisfied as long as $|\theta_h|+|\theta_w|<1$.[28]

The reduced form of (A.3) is obtained replacing (A.4) (for $k=1$, which results in a first-order "approximate" inverse) into (A.3):

$$\Delta\mathbf{c}\approx(\mathbf{S}'_h\mathbf{G}\gamma_h+\mathbf{S}'_w\mathbf{G}\gamma_w+\mathbf{S}'_h\delta_h+\mathbf{S}'_w\delta_w)\Delta\mathbf{X}$$
$$+(\theta_h\mathbf{G}_h+\theta_w\mathbf{G}_w)(\mathbf{S}'_h\mathbf{G}\gamma_h+\mathbf{S}'_w\mathbf{G}\gamma_w+\mathbf{S}'_h\delta_h+\mathbf{S}'_w\delta_w)\Delta\mathbf{X}+\mathbf{v}.$$

The interesting part of the identification argument, which is informal due to the approximation of (A.4) is that one derives identification power from the cross-products between the different G matrices (in the case considered by Bramoullé *et al.*, 2009, the population is made of single individuals, hence identification comes only from powers of the adjacency matrix; but see Blume *et al.*, 2015, for an idea similar to the one we are using here). In the equation above, all the parameters of interest are separately identified as long as S'_hG, S'_wG, S'_h, S'_w, $G_hS'_hG$, $G_wS'_wG$, $G_hS'_h$, and $G_wS'_w$ are linearly independent.

B. VARIABLES DEFINITION

In our regressions, we control for region of residence dummies, industry dummies, and firm legal type dummies, categorized as follows:

Region of residence: Copenhagen, Broader Copenhagen, Frederiksborg, Roskilde, Vestsjælland, Storstrøm, Bornholm, Fyn, Sønderjylland, Ribe, Vejle, Viborg, Aarhus, Ringkøbing, and Nordjylland.

Industry: Agriculture, fishing and mining; manufacturing; utilities; construction; retail trade, hotels and resturants; transportation, storage and communication; financial intermediation and business; public and personnal services; Other.

Firm legal type: Limited liability ApS (Ltd.), Publicly traded limited liability A/S (Inc.), and other.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

Acknowledgments. We thank the Editor, four anonymous referees, Orazio Attanasio, Chris Carroll, Nicola Fuchs-Schuendeln, Soren Leth-Petersen, Magne Mogstad, Petra Persson, Tomas Rodriguez Barraquer, Nick Roussanov, Dan Silverman, Nick Souleles, Adam Szeidl, and seminar participants at several seminars and conferences in Europe and the U.S. for comments. Financial support from the Spanish Ministry of Economy and Competitiveness through the Severo Ochoa Programme for Centres of Excellence in R&D (SEV-2011-0075) and ECO2011-28822, ECO2014-54401-P, the EU through the Marie Curie CIG grant FP7-631510 the Swiss National Science Foundation grant 100018_182243 (De Giorgi), the ERC starting grant 284024 and the NSF grant 1458536 (Pistaferri) is gratefully acknowledged.

REFERENCES

ABOWD, J. M., KRAMARZ, F. and MARGOLIS, D. N. (1999), "High Wage Workers and High Wage Firms", *Econometrica*, **67**, 251–333.

AGARWAL, S., QIAN, W. and ZOU, X. (2017), "Thy Neighbor's Misfortune: Peer Effect on Consumption" (Mimeo).

BANERJEE, A., CHANDRASEKHAR, A., DUFLO, E., *ET AL.* (2013). "The Diffusion of Microfinance", *Science*, **341**, 1236498.

28. The sufficient condition for ensuring $(\mathbf{I}-(\theta_h\mathbf{G}_h+\theta_w\mathbf{G}_w))^{-1}=\sum_{k=0}^{\infty}(\theta_h\mathbf{G}_h+\theta_w\mathbf{G}_w)^k$ is that $\|\theta_h\mathbf{G}_h+\theta_w\mathbf{G}_w\|_n<1$ for some operator matrix norm n (see Meyer, 2000, p. 527). This is useful because we can use the fact that $\|\mathbf{G}_h\|_\infty=\|\mathbf{G}_w\|_\infty=1$ due to the row-normalization of the adjacency matrix. Moreover, we can use the following two properties of matrix norms: (a) $\|\alpha\mathbf{A}\|=|\alpha|\,\|\mathbf{A}\|$, and (b) $\|\mathbf{A}+\mathbf{B}\|\le\|\mathbf{A}\|+\|\mathbf{B}\|$, if α is a scalar and $\mathbf{A}$ and $\mathbf{B}$ are square matrices. Hence: $\|\theta_h\mathbf{G}_h+\theta_w\mathbf{G}_w\|_\infty\le\|\theta_h\mathbf{G}_h\|_\infty+\|\theta_w\mathbf{G}_w\|_\infty=|\theta_h|\,\|\mathbf{G}_h\|_\infty+|\theta_w|\,\|\mathbf{G}_w\|_\infty=|\theta_h|+|\theta_w|$. Hence, the condition $\|\theta_h\mathbf{G}_h+\theta_w\mathbf{G}_w\|_\infty<1$ is clearly satisfied if $|\theta_h|+|\theta_w|<1$.

BERTRAND, M., LUTTMER, E. F. P. and MULLAINATHAN, S. (2000). "Network Effects and Welfare Cultures," *Quarterly Journal of Economics*, **115**, 1019–1055.

BLUME, L., BROCK, W., DURLAUF, S., *ET AL.* (2015), "Linear Social Interactions Models", *Journal of Political Economy*, **123**, 444–496.

BLUNDELL, R., BROWNING, M. and MEGHIR, C. (1994). "Consumer Demand and the Life-Cycle Allocation of Household Expenditures", *The Review of Economic Studies*, **61**, 57–80.

BRAMOULLÉ, Y., DJEBBARI, H. and FORTIN, B. (2009). "Identification of Peer Effects Through Social Networks", *Journal of Econometrics*, **150**, 41–55.

BROCK, W. and DURLAUF, S. (2001). "Interaction-Based Models", in Heckman J. and Leamer E. (eds) *Handbook of Econometrics*, Vol. 5 (Amsterdam: North-Holland).

BROWNING, M., DEATON, A. and IRISH, M. (1985). "A Profitable Approach to Labor Supply and Commodity Demands Over the Life Cycle", *Econometrica*, **53**, 503–543.

BROWNING, M. and LETH-PETERSEN, S. (2003). "Imputing Consumption from Income and Wealth Information", *The Economic Journal*, **113**, F282–F301.

BROWNING, M. and MEGHIR, C. (1991). "The Effects of Male and Female Labor Supply on Commodity Demands", *Econometrica*, **59**, 925–951.

CALVÓ-ARMENGOL, A., PATACCHINI, E. and ZENOU, Y. (2009). "Peer Effects and Social Networks in Education", *Review of Economic Studies*, **76**, 1239–1267.

CARD, D., DEVICIENTI, F. and MAIDA, A. (2014), "Rent-Sharing, Holdup, and Wages: Evidence from Matched Panel Data", *Review of Economic Studies*, **81**, 84–111.

CARROLL, C. D., OVERLAND, J. and WEIL, D. N. (1997). "Comparison Utility in a Growth Model", *Journal of economic growth*, **2**, 339–367.

CAEYERS, B. and FAFCHAMPS, M. (2016). "Exclusion Bias in the Estimation of Peer Effects" (NBER Working Paper 22565).

CHARLES, K., HURST, E. and ROUSSANOV, N. (2009). "Conspicuous Consumption and Race", *Quarterly Journal of Economics*, **124**, 425–467.

CHETTY, R., FRIEDMAN, J. N., OLSEN, T., *ET AL.* (2011). "Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records", *The Quarterly Journal of Economics*, **126**, 749–804.

CURRARINI, S., JACKSON, M. and PIN, P. (2009). "An Economic Model of Friendship: Homophily, Minorities, and Segregation", *Econometrica*, **77**, 1003–1045.

DEATON, A. and MUELLBAUER, J. (1980). "An almost ideal demand system", *The American Economic Review*, **70**, 312–326.

DE GIORGI, G., PELLIZZARI, M. and REDAELLI, S. (2010). "Identification of Peer Effects Through Partially Overlapping Peer Groups", *The American Economic Journal: Applied Economics*, **2**, 241–275.

DE GIORGI, G., FREDERIKSEN, A. and PISTAFERRI, L. (2016), "Consumption Network Effects" (Working Paper No. 22357, NBER).

DE PAULA, A., RASUL, I. and SOUZA, P. (2018), "Recovering Social Networks from Panel Data: Identification, Simulations and an Application" (CEPR DP 12792).

DE PAULA, A. (2017), "Econometrics of Network Models", in Honore, B., Pakes, A., Piazzesi, M., *et al.* (eds) *Advances in Economics and Econometrics: Theory and Applications: Eleventh World Congress (Econometric Society Monographs)* (Cambridge: Cambridge University Press) 268–323.

DUESENBERRY, J. S. (1948). "Income-Consumption Relations and Their Implications", in Metzler, L. and A. H. Hansen (eds) *Income, Employment and Public Policy*: Essays in honour of Alvin H. Hansen (New York: WW Norton & Company, Inc.).

DUFLO, E. and SAEZ, E. (2003). "The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment", *Quarterly Journal of Economics*, **118**, 815–842.

FAGERENG, A., GUISO, L., MALACRINO, D., *ET AL.* (2016). "Heterogeneity and Persistence in Returns to Wealth" (Working Paper No. 22822, NBER).

FAGERENG, A., GUISO, L. and PISTAFERRI, L. (2018)."Portfolio Choices, Firm Shocks and Uninsurable Wage Risk", *Review of Economic Studies*, **85**, 396–436.

GALI, J. (1994). "Keeping Up with the Joneses: Consumption Externalities, Portfolio Choice, and Asset Prices", *Journal of Money, Credit and Banking*, **26**, 1–8.

GLAESER, E. L., SCHEINKMAN, J. A. and SACERDOTE, B. (2003). "The Social Multiplier", *Journal of the European Economic Association P&P*, 1(2/3), 345–353.

GOLDSMITH-PINKHAM, P. and IMBENS, G. W. (2013). "Social Networks and the Identification of Peer Effects", *Journal of Business & Economic Statistics*, **31**, 253–264.

GUISO, L., PISTAFERRI, L. and SCHIVARDI, F. (2005). "Insurance within the firm", *Journal of Political Economy*, **113**, 1054–1087.

HEFFETZ, O. (2011). "A Test of Conspicuous Consumption: Visibility and Income Elasticities", *Review of Economics and Statistics*, **93**, 1101–1117.

HOLZER, H.J. (1988). "Search Method Use by Unemployed Youth", *Journal of Labor Economics*, **6**, 1–20.

JAPPELLI, T. and PISTAFERRI, L. (2014). "Fiscal Policy and MPC Heterogeneity", *American Economic Journal: Macroeconomics*, **6**, 107–136.

KLINE, P., PETKOVA, N., WILLIAMS, H., *ET AL.* (2018), "Who Profits from Patents? Rent Sharing at Innovative Firms" NBER Working Paper 25245.

KUHN, P., KOOREMAN, P., P., SOETEVENT, A., *et al.* (2011). "The Effects of Lottery Prizes on Winners and Their Neighbors: Evidence from the Dutch Postcode Lottery", *The American Economic Review*, **101**, 2226–2247.

LEWBEL, A., NORRIS, S., PENDAKUR, K. and QU, X. (2018), "Consumption Peer Effects and Utility Needs in India", unpublished manuscript.

LJUNGQVIST, L. and UHLIG, H. (2000). "Tax Policy and Aggregate Demand Management Under Catching Up with the Joneses", *American Economic Review*, **90**, 356–366.

MACE, B. J. (1991). "Full Insurance in the Presence of Aggregate Uncertainty", *Journal of Political Economy*, **99**, 928–956.

MANSKI, C. (1993). "Identification of Endogenous Social Effects: The Reflection Problem", *Review of Economic Studies*, **60**, 531–542.

MAS, A. and MORETTI, E. (2009). "Peers at Work", *American Economic Review*, **99**, 112–145.

MAURER, J. and MEIER, A. (2008). "Smooth It Like the 'Joneses'? Estimating Peer-Group Effects in Intertemporal Consumption Choice", *The Economic Journal*, **118**, 454–476.

MEYER, C. (2000). *Matrix Analysis and Applied Linear Algebra.* (Philadelphia: Society for Industrial and Applied Mathematics).

MOFFITT, R. (2001). "Policy Interventions, Low-Level Equilibria, and Social Interactions", in Durlauf S. and Young, H. P. (eds) *Social Dynamics* (Cambridge: MIT Press).

MONTGOMERY, J. D. (1991). "Social Networks and Labor-market Outcomes: Toward an Economic Analysis". *The American Economic Review*, **81**, 1408–1418.

PISTAFERRI, L. (1999). "Informal Networks in the Italian Labor Market", *Giornale degli economisti e annali di economia*, **58**, 355–375.

SACERDOTE, B. (2001). "Peer Effects with Random Assignment: Results for Dartmouth Roommates", *Quarterly Journal of Economics*, **116**, 681–704.

SANDERSON, E. and WINDMEIJER, F. (2016). "A Weak Instrument F-Test in Linear IV Models with Multiple Endogenous Variables", *Journal of Econometrics*, **190**, 212–221.

VEBLEN, T. (1899). *The Theory of the Leisure Class: An Economic Study in the Evolution of Institutions* (London: Macmillan).