

Lewbel, Arthur; Norris, Samuel; Pendakur, Krishna; Qu, Xi

Article

Consumption peer effects and utility needs in India

Quantitative Economics

Provided in Cooperation with:

The Econometric Society

Suggested Citation: Lewbel, Arthur; Norris, Samuel; Pendakur, Krishna; Qu, Xi (2022) : Consumption peer effects and utility needs in India, Quantitative Economics, ISSN 1759-7331, The Econometric Society, New Haven, CT, Vol. 13, Iss. 3, pp. 1257-1295, <https://doi.org/10.3982/QE1760>

This Version is available at:

<https://hdl.handle.net/10419/296300>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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



<https://creativecommons.org/licenses/by-nc/4.0/>

Consumption peer effects and utility needs in India

ARTHUR LEWBEL

Department of Economics, Boston College

SAMUEL NORRIS

Department of Economics, University of British Columbia

KRISHNA PENDAKUR

Department of Economics, Simon Fraser University

XI QU

Antai College of Economics and Management, Shanghai Jiao Tong University

We construct a peer effects model where mean expenditures of consumers in one's peer group affect utility through perceived consumption needs. We provide a novel method for obtaining identification in social interactions models like ours, using ordinary survey data, where very few members of each peer group are observed. We implement the model using standard household-level consumer expenditure survey microdata from India. We find that each additional rupee spent by one's peers increases perceived needs, and thereby reduces one's utility, by the equivalent of a 0.25 rupee decrease in one's own expenditures. These peer costs may be larger for richer households, meaning transfers from rich to poor could improve even inequality-neutral social welfare, by reducing peer consumption externalities. We show welfare gains of billions of dollars per year might be possible by replacing government transfers of private goods to households with providing public goods or services, to reduce peer effects.

KEYWORDS. Consumer demand, consumption, measurement error, welfare, peer effects.

JEL CLASSIFICATION. C21, C31, D12.

1. INTRODUCTION

There are substantial peer effects in income and consumption. People's evaluation of their own income depends on the income of their peers (Kahneman (1992), Luttmer

Arthur Lewbel: lewbel@bc.edu

Samuel Norris: sam.norris@ubc.ca

Krishna Pendakur: pendakur@sfu.ca

Xi Qu: xiqu@sjtu.edu.cn

We thank Chris Muris and Niharika Singh for comments on this work. We also thank seminar participants at Harvard, Boston University, University of Pennsylvania, Vanderbilt, and Rice. Norris and Pendakur are grateful for financial support from the Social Sciences and Humanities Research Council of Canada through its Doctoral Fellowship Awards and Insight Research Grants Programs, respectively.

© 2022 The Authors. Licensed under the [Creative Commons Attribution-NonCommercial License 4.0](https://creativecommons.org/licenses/by-nc/4.0/). Available at <http://qeconomics.org>. <https://doi.org/10.3982/QE1760>

(2005), and Clark, Frijters, and Shields (2008)). Their consumption choices also depend on those of their peers (Gali (1994), Boneva (2013), and de Giorgi, Frederiksen, and Pistaferri (2020)), and the perceived value of individual goods or brands depends on the consumption of those goods in relevant reference groups (Rabin (1998), Kalyanaram and Winer (1995), and Chao and Schor (1998)).

Despite the strong evidence of peer effects in consumption choices, there has been much less work evaluating the resulting welfare impacts of these effects. In this paper, we study the impact of changes in peer mean expenditures on utility, asking how much one's own expenditure would have to increase to compensate for a unit increase in peer group expenditures.

One way to measure the welfare impacts of peer effects would be to directly regress an observed utility measure (i.e., stated well-being) on own and peer expenditures, as in Luttmer (2005). This has the drawback of relying on coarse self-reports of well-being, which generally suffer from lack of interpersonal comparability, framing biases, measurement errors, and problems of interpretation. Most empirical studies of consumption peer effects instead regress individual consumption on mean peer group consumption and other covariates (Chao and Schor (1998) and Boneva (2013)). Such regressions reveal behavioral responses to peer expenditures, but do not reveal the welfare implications of these peer effects.

To study the welfare effects of peer consumption, we propose a “keeping up with the Joneses”-type structural model that exploits revealed preference based links between consumption and utility to recover the welfare implications of peer expenditures on consumption behavior. In our model, one's perceived required expenditures, or “needs,” depend on, among other things, the mean expenditures of one's peer group. The higher are these perceived needs, the more one must spend to attain the same level of utility. Consistent with other empirical evidence, we find that consumers lose utility from feeling poorer when their peers get richer, and so consumers feel they must spend and consume more when their peers consume more (Luttmer (2005), Dynan and Ravina (2007), and Clark and Senik (2010)). In contrast to those papers, which use direct data on reported well-being, we identify the money-metric costs of peer consumption from ordinary consumer demand data.

In estimating peer effects, much progress has been made in overcoming the endogeneity of peer data, and allowing for fixed or random peer group effects, by using detailed social network information. For example, de Giorgi, Frederiksen, and Pistaferri (2020) instrument for peer consumption with friend-of-friend consumption. However, in our application we use only standard cross-section expenditure survey data, of the type that is commonly collected by governments all over the world. As a result, we cannot make use of network information or variation in peer group sizes (as in Lee (2007)) to obtain identification.

We estimate our model using consumption survey data from India.¹ Our groups are defined at a very local geographic level; roughly a small neighborhood. Within neighborhoods, we also group people by religion and caste. This results in groups comprised

¹A number of studies find significant peer effects in India, for example, Banerjee, Chandrasekhar, Duflo, and Jackson (2013).

of a few hundred members, but we only observe ten or fewer of the members in each group. This gives rise to some unusual econometric obstacles that must be overcome.

One such obstacle is that we cannot consistently estimate within-group mean expenditures, because so few members of each group are observed. Another is identifying the impact of group level peer effects while allowing for unobserved heterogeneity in group behavior, in the form of group level fixed effects or random effects. A third is coping with nonlinearities associated with maximization of empirically plausible utility functions. We propose some novel identifying moments to obtain model identification despite these obstacles, and provide an associated GMM estimator. These innovations in the econometric identification of peer effects could prove useful in other applications of social interactions models.

Empirically, we find that dealing with the peer group mean measurement error issue is particularly important. Failing to account for these errors leads to attenuation biases so large that the estimated peer effects are reduced by up to 90% in some specifications. [Chandrasekhar and Lewis \(2011\)](#) also document mean biases of 90% due to incomplete measurement of some networks.

Our empirical results show that an increase in spending by one's peers of one rupee has the same effect on one's utility as a decrease in one's own expenditures of about one-fourth of a rupee. We also find some evidence that peer effects are smaller for lower socioeconomic status groups.

These results have important implications for tax and redistribution policy. First, they suggest that personal taxes may be less costly in terms of social welfare than is implied by standard demand models, which ignore peer effects. Since peer effects are a negative consumption externality, the reduction in consumption caused by taxes does not reduce welfare as much it would in the absence of that externality.

Second, if the utility associated with public goods or government services are not subject to these peer effects (or engender smaller peer effects), then governments can increase welfare by substituting the provision of public goods for the provision of private goods. This effect can be very large: we perform a rough calculation, which shows that replacing India's food subsidy program with more generous provision of public goods and services, such as public sanitation or cleaner air, could potentially increase money metric welfare by up to 180 billion rupees (2.5 billion US dollars) per year at no additional cost, by reducing peer effects.

Third, the finding that poorer households may have smaller peer effects suggests that transfers from higher to lower status groups can increase total welfare, by reducing peer externalities. The usual argument for transfers of money from rich to poor is the assumption of declining social marginal welfare (meaning society benefits more by giving the poor an additional dollar than it loses by taking away a dollar from the rich, *ceteris paribus*). In contrast, our results give a reason why progressive taxation can increase social welfare, even if all consumers have the same marginal utility of money, and even if one's social welfare function is inequality-neutral.

Consider a model where each consumer, indexed by i , is a member of a peer group, indexed by g . Let \mathbf{q}_i be the vector of (continuous) quantities of goods that consumer i consumes. We specify and estimate utility derived demand functions, which express \mathbf{q}_i as a function of prices, a total expenditure budget x_i , and demographic characteristics. Consumer i is a member of a peer group g . Let $\bar{\mathbf{q}}_g$ be the expected value of the quantity vector \mathbf{q} among all the members of group g .² The peer effects in our model have demand functions that also depend nonlinearly on $\bar{\mathbf{q}}_g$. There is a long history, going back to Samuelson (1947), of modeling needs in utility and demand functions as analogous to fixed costs or overheads in production. In our model, these needs depend on $\bar{\mathbf{q}}_g$, and \mathbf{q}_i in turn depends on needs.

Given estimates of our demand functions, and hence (by revealed preference theory) ordinal utility functions, we can answer the question: If peer spending $\bar{\mathbf{q}}_g$ increases, how much poorer does consumer i feel? More precisely, how much more would consumer i need to spend (i.e., how much would his or her budget x_i need to increase) to give that consumer the same level of utility she had before $\bar{\mathbf{q}}_g$ increased? This is the fundamental welfare question our analysis seeks to answer.

Our system of demand functions is an example of a social interactions model, since it includes the vector of group means $\bar{\mathbf{q}}_g$ as covariates. Our model differs from standard social interactions models (e.g., Manski (1993, 2000), Brock and Durlauf (2001), Lee (2007), and Blume et al. (2011)) in a variety of ways. First, our model is nonlinear and vector-valued while most such models are linear and scalar-valued. This nonlinearity complicates some aspects of identification, but it helps overcome other issues, such as allowing us to include group-level fixed effects in the model (in a linear model, differencing or demeaning to remove fixed effects would also remove $\bar{\mathbf{q}}_g$).

A second difference is that most social interactions models make use of network information for identification, but we cannot. Examples of such network information include the use of exogenous variation in group composition or size (e.g., Lee (2007); Carrell, Fullerton, and West (2009) and Duflo, Dupas, and Kremer (2011)), or the use of detailed network structure like intransitive triads, where data on friends of friends provides instruments for identification (e.g., Bramoullé, Djebbari, and Fortin (2009), De Giorgi, Pellizzari, and Redaelli (2010), Jochmans and Weidner (2016), or de Giorgi, Frederiksen, and Pistaferri (2020)). In contrast, our model uses consumer expenditure survey data of the type that many countries collect for constructing consumer price indices. Since such surveys do not contain social network information, we can only define peer groups based on demographic characteristics and geography, and we therefore cannot exploit any network structure to help identification.

A third difference from most social interactions applications is that, having survey data, we only observe a small number of the members in each peer group. Our peer groups each have a few hundred members or more in the population, but we observe

²Letting $\bar{\mathbf{q}}_g$ be the expected value rather than the average of \mathbf{q}_i among all group members mitigates some technical issues, such as the reflection problem, and it means members do not need to consider the effect of their own decisions on $\bar{\mathbf{q}}_g$.

at most 10 members of each group in our sample. As a result, we cannot consistently estimate group means $\bar{\mathbf{q}}_g$. For each group g , we can at best construct an estimate $\hat{\mathbf{q}}_g$ by averaging across the small number of members that we do observe in each group. This greatly complicates identification and estimation of our model, because replacing $\bar{\mathbf{q}}_g$ with $\hat{\mathbf{q}}_g$ introduces group level measurement error into the model, and this measurement error $\hat{\mathbf{q}}_g - \bar{\mathbf{q}}_g$ is both endogenous and correlated with other components of the model. These problems are further exacerbated by nonlinearity of the demand function, resulting in interaction terms like the measurement error $\hat{\mathbf{q}}_g - \bar{\mathbf{q}}_g$ multiplied by x_i (the budget of consumer i).

Section 3 below illustrates our general procedure for dealing with these econometric issues, in the context of a simple generic quadratic model. This procedure should be of independent interest to others wishing to estimate peer effects using survey data. In some Appendices, we provide a sequence of theorems proving that our identification method and associated proposed estimators yield consistent estimators and valid inference, both for the simple quadratic generic model (with some extensions), and for our general utility derived demand model.

The remainder of this paper proceeds as follows. In Section 2, we expand on the structural model of utility, demand and peer effects introduced above. Section 3 summarizes our general results on identification and estimation. Our empirical estimates are presented in Section 4, with policy implications provided in Section 5. Section 6 then concludes.

2. UTILITY AND DEMAND WITH PEER EFFECTS IN NEEDS

There is a long literature that connects utility and well-being to peer income or consumption levels (see, e.g., Frank (1999, 2012)). The Easterlin (1974) paradox asserts an empirical connection between well-being and national average incomes. Though the strength of this connection is debated (Stevenson and Wolfers (2008)), the correlation between utility and national-level consumption, *ceteris paribus*, appears negative. Dynan and Ravina (2007) and Clark and Senik (2010) regress self-reported utility on own budgets and national average budgets, and other correlated aggregate measures like inequality, and find that this negative correlation still stands. Similar results hold for much smaller reference groups, for example, Luttmer (2005) finds that an increase of the average income in one's neighbors reduces self-reported well being.

The possible mechanisms for this correlation are varied. Veblen (1899) effects make consumers value consumption of visible status goods. Reference-dependent utility functions hinge preferences on own-endowments (Kahneman and Tversky (1979)). More recent work on these models has led to reference-dependence that is "other-regarding," where utilities depend on reference points that are driven by other agents' decisions or endowments. Models of "keeping up with the Joneses" have one's own consumption feel smaller when one's peers consume more. Surveys of this literature include Kahneman (1992) and Clark, Frijters, and Shields (2008).

Taken together, this literature suggests that the utility of consumer i should depend on both \mathbf{q}_i and $\bar{\mathbf{q}}_g$, and that utility is increasing in \mathbf{q}_i and decreasing in $\bar{\mathbf{q}}_g$.^{3,4,5} If we could observe utility and consumption quantities of individuals and groups, we could directly test this. Luttmer (2005) estimates an approximation of this relationship, by regressing a crude measure of utility (reported life satisfaction on a coarse ordinal scale) not on \mathbf{q}_i and $\bar{\mathbf{q}}_g$, but on x_i and its group mean \bar{x}_g . Separate from our main empirical application, we estimate a similar regression, using data from India and groups that are roughly comparable to those in our main empirical analysis. The results agree with Luttmer (2005) and support our main model's underlying assumption that increases in peer expenditures decrease rather than increase utility. Our main model does not depend on crude utility measures, but instead identifies comparable structural parameters obtained from utility-derived demand functions via revealed preference.

A number of papers relate consumption choices to peer consumption levels, although these analyses are essentially nonstructural (Chao and Schor (1998), Boneva (2013), and de Giorgi, Frederiksen, and Pistaferri (2020)). All these papers suggest that the magnitudes of peer effects in consumption choices are large. In our notation, these papers use empirical approaches analogous to regressing \mathbf{q}_i on x_i and $\bar{\mathbf{q}}_g$. However, establishing how much consumption \mathbf{q}_i changes when peer consumption $\bar{\mathbf{q}}_g$ changes does not answer the welfare question of how $\bar{\mathbf{q}}_g$ affects utility, and hence how much one would need to increase x_i to compensate for the loss of utility from an increase in $\bar{\mathbf{q}}_g$. Answering this type of welfare question requires linking expenditures to utility, which is what our structural model does.

2.1 The utility-derived demand model

Our model is that each consumer, indexed by i , is a member of a peer group, indexed by g . Note that g should have a subscript i , denoting the particular group that contains consumer i , but we drop this subscript to avoid notational clutter. Let \mathbf{q}_i be the vector of (continuous) quantities of goods that consumer i consumes. Utility is given by $U_i = U(\mathbf{q}_i - \mathbf{f}_i)$, where U_i is the attained utility level of consumer i , U is a utility function (ignoring taste heterogeneity for now), and \mathbf{f}_i is a vector of the needs of consumer i .

Needs \mathbf{f}_i is a quantity vector, with elements equal to the minimum quantities that consumer i must consume of each good before he or she starts to get any utility. In the context of a linear model, Samuelson (1947) defines the quantity vector \mathbf{f}_i as the “necessary set” of goods. The Stone (1954) and Geary (1949) linear expenditure system is just

³It is of course possible that peer group expenditures matter in other ways than just through group means $\bar{\mathbf{q}}_g$. We only consider group means here because of data limitations and other econometric issues discussed later.

⁴One could imagine utility positively correlated with $\bar{\mathbf{q}}_g$, for example, through happiness for the success of one's peers. But the empirical evidence, including our own results, suggest that the correlation is negative.

⁵Our groups are defined (in the main) by geography. This implies a substantial risk of misspecifying how consumers are assigned to peer groups. We mitigate this risk in part by constructing very small groups, since defining groups that are too small creates inefficiency but not bias. We also show how misspecification of groups will generally lead to downward bias in peer effects estimates, so our estimated effects are likely to be conservative, and we perform some specification tests regarding group definitions.

a Cobb–Douglas utility function U with needs equalling a constant vector \mathbf{f} . Gorman (1976) analyzed the general form $U_i = U(\mathbf{q}_i - \mathbf{f}_i)$ for arbitrary utility functions U , letting \mathbf{f}_i depend on demographic variables and taste shifters \mathbf{z}_i .

Let overbars indicate true within-group means, and hats indicate sample averages. We extend Gorman (1976) by letting needs \mathbf{f}_i also depend on $\bar{\mathbf{q}}_g$, the expected value of the quantity vector \mathbf{q} among the members of consumer i 's peer group g . The model therefore has $\mathbf{f}_i = \mathbf{f}(\mathbf{z}_i, \bar{\mathbf{q}}_g)$ for a vector valued needs function \mathbf{f} . Let \mathbf{p} be the price vector corresponding to \mathbf{q}_i , and let x_i be consumer i 's budget (total expenditures). Consumer i chooses the vector \mathbf{q}_i to maximize his or her utility

$$U_i = U(\mathbf{q}_i - \mathbf{f}(\mathbf{z}_i, \bar{\mathbf{q}}_g)) \quad (1)$$

under the linear budget constraint $\mathbf{p}'\mathbf{q}_i \leq x_i$.

One can equivalently represent preferences using an indirect utility function, defined as the maximum utility attainable with a given budget x_i when facing prices \mathbf{p} . Gorman (1976) shows⁶ that for any regular utility function in this form, there exists a corresponding indirect utility function V such that

$$U_i = V(\mathbf{p}, x_i - \mathbf{p}'\mathbf{f}(\mathbf{z}_i, \bar{\mathbf{q}}_g)). \quad (2)$$

Indirect utility functions of this form can be shown to have many desirable properties for welfare calculations.⁷ Blackorby and Donaldson (1994) and Donaldson and Pendakur (2006) show that the function \mathbf{f} (without $\bar{\mathbf{q}}_g$) is uniquely identified up to location from consumer demand functions. We show later that we can also uniquely identify how \mathbf{f} depends on $\bar{\mathbf{q}}_g$.

Luttmer (2005) regresses a self-reported measure of happiness on \mathbf{z}_i , y_i , and \hat{y}_g (where for Luttmer, y_i is the income of consumer i , and \hat{y}_g is the observed within-group average income). We can interpret his regression as a simplified and linearized version of equation (2), where self-reported happiness is assumed to proxy for U_i , income y_i replaces x_i , and all the effects of $\bar{\mathbf{q}}_g$ are subsumed by \hat{y}_g . Table 1 (column 3) in Luttmer (2005) gives endogeneity-corrected estimates of the coefficients of \hat{y}_g and y_i of -0.296 and 0.361 , respectively. The negative ratio of these is 0.82 , meaning that a 100 dollar increase in group-average income has the same effect on reported happiness as an 82 dollar reduction in own-income. We later estimate an object that has a comparable interpretation to this relative coefficient. But instead of assuming that U_i equals an observed happiness measure that can be compared across individuals and regressed on covariates, we let U_i be unobserved. We instead derive demand equations from equation (2), and then recover the implied peer effects on utility.

The demand functions that result from maximizing our utility function can be obtained by applying Roy's (1947) identity to the indirect utility function of equation (2).

⁶His version did not include $\bar{\mathbf{q}}_g$.

⁷Blackorby and Donaldson (1994) show that indirect utility functions $U_i = V(\mathbf{p}, x_i - \mathbf{p}'\mathbf{f}_i)$ satisfy Absolute Equivalence Scale Exactness (AESE). For preferences that satisfy AESE, one can define equivalent income as $x_i - \mathbf{p}'\mathbf{f}_i$ and show that the sum of equivalent income across consumers is a valid money-metric based social welfare function.

These demand functions have the form $\mathbf{q}_i = \mathbf{h}(\mathbf{p}, x - \mathbf{p}'\mathbf{f}_i) + \mathbf{f}_i$, where $\mathbf{f}_i = \mathbf{f}(\mathbf{z}_i, \bar{\mathbf{q}}_g)$ and the vector valued function \mathbf{h} is defined by $\mathbf{h}(\mathbf{p}, x) = -\nabla_{\mathbf{p}}V(\mathbf{p}, x)/\nabla_x V(\mathbf{p}, x)$. See, for example, Pollak and Wales (1981) and Pendakur (2005).

To allow for unobserved heterogeneity across consumers, we append the error term $\mathbf{v}_g + \mathbf{u}_i$ to these demand functions, so

$$\mathbf{q}_i = \mathbf{h}(\mathbf{p}, x_i - \mathbf{p}'\mathbf{f}(\mathbf{z}_i, \bar{\mathbf{q}}_g)) + \mathbf{f}(\mathbf{z}_i, \bar{\mathbf{q}}_g) + \mathbf{v}_g + \mathbf{u}_i, \quad (3)$$

where \mathbf{v}_g is a J -vector of group level fixed or random effects and \mathbf{u}_i is a J -vector of individual specific error terms that are assumed to have zero means conditional on x_i , \mathbf{z}_i , and \mathbf{p} . This model can alternatively be interpreted as including $\mathbf{v}_g + \mathbf{u}_i$ additively in \mathbf{f}_i but imposing the (somewhat peculiar) restriction that $\mathbf{p}'(\mathbf{v}_g + \mathbf{u}_i) = 0$. If $\mathbf{v}_g + \mathbf{u}_i$ is a component of needs $\mathbf{v}_g + \mathbf{u}_i$, then the restriction $\mathbf{p}'(\mathbf{v}_g + \mathbf{u}_i) = 0$ would be needed to keep each individual on their budget constraint, and it makes these errors drop out of \mathbf{h} .

The difference between interpreting $\mathbf{v}_g + \mathbf{u}_i$ as a departure from utility maximization or as unobserved preference heterogeneity is irrelevant for identification and estimation. However, the difference can affect whether it is appropriate to include $\mathbf{v}_g + \mathbf{u}_i$ in welfare calculations or not. However, all of the welfare analyses we perform are based on changes in utilities and in fixed costs, rather than levels, so these terms, if they were included in our welfare measures, would just get differenced out anyway. Note that we will control for, but not actually estimate, the fixed or random effects \mathbf{v}_g .

In the fixed-effects model, \mathbf{v}_g can be correlated in unknown ways with regressors including \mathbf{p} , x , and $\bar{\mathbf{q}}_g$. The random-effects model imposes the additional restriction that \mathbf{v}_g be independent of regressors. As a result, the random-effects model will be much more efficient, but at the cost of imposing these possibly questionable independence restrictions.

We take the needs function $\mathbf{f}(\mathbf{z}_i, \bar{\mathbf{q}}_g)$ to be linear, so

$$\mathbf{f}_i = \mathbf{A}\bar{\mathbf{q}}_g + \mathbf{C}\mathbf{z}_i \quad (4)$$

for some matrices of parameters \mathbf{A} and \mathbf{C} . Linearity of \mathbf{f}_i in \mathbf{z}_i is commonly assumed in empirical demand analysis, so we extend that linearity to the additional variables $\bar{\mathbf{q}}_g$. The vector of demand functions given by equation (4) then reduce to

$$\mathbf{q}_i = \mathbf{h}(\mathbf{p}, x_i - \mathbf{p}'\mathbf{A}\bar{\mathbf{q}}_g - \mathbf{p}'\mathbf{C}\mathbf{z}_i) + \mathbf{A}\bar{\mathbf{q}}_g + \mathbf{C}\mathbf{z}_i + \mathbf{v}_g + \mathbf{u}_i. \quad (5)$$

2.2 Utility and demand functions

To obtain equations we will estimate, we need to specify the indirect utility function V , which then determines the vector-valued function \mathbf{h} . Based on a long empirical literature,⁸ we assume

$$V(\mathbf{p}, x) = -(x - R(\mathbf{p}) - \mathbf{p}'(\mathbf{A}\bar{\mathbf{q}}_g + \mathbf{C}\mathbf{z}_i))^{-1}B(\mathbf{p}) - D(\mathbf{p}) \quad (6)$$

⁸Many studies of commodity demands have found that observed demand functions are close to the polynomial. See, for example, Lewbel (1991), Banks, Blundell, and Lewbel (1997), and references therein. Gorman (1981) shows that any polynomial demand system has a maximum rank of three. Lewbel (1989) provides the tractable classes of indirect utility functions that yield rank three polynomials. The most com-

for some differentiable functions R , B , and D . Applying Roy's identity to obtain the function \mathbf{h} and equation (5) yields demand equations

$$\mathbf{q}_i = (x_i - R(\mathbf{p}) - \mathbf{p}'(\mathbf{A}\bar{\mathbf{q}}_g + \mathbf{C}\mathbf{z}_i))^2 \frac{\nabla D(\mathbf{p})}{B(\mathbf{p})} + (x_i - R(\mathbf{p}) - \mathbf{p}'(\mathbf{A}\bar{\mathbf{q}}_g + \mathbf{C}\mathbf{z}_i)) \frac{\nabla B(\mathbf{p})}{B(\mathbf{p})} + \nabla R(\mathbf{p}) + \mathbf{A}\bar{\mathbf{q}}_g + \mathbf{C}\mathbf{z}_i + \mathbf{v}_g + \mathbf{u}_i. \quad (7)$$

Rationality (consistency with utility maximization) requires that $R(\mathbf{p})$ and $B(\mathbf{p})$ be homogeneous of degree 1 in \mathbf{p} and that $D(\mathbf{p})$ be homogeneous of degree 0 in \mathbf{p} . Standard functions that satisfy these conditions and yield price-flexible (in the sense of Diewert (1974)) demand functions are $R(\mathbf{p}) = \mathbf{p}^{1/2'} \mathbf{R} \mathbf{p}^{1/2}$ where \mathbf{R} is a symmetric matrix, $\ln B(\mathbf{p}) = \mathbf{b}' \ln \mathbf{p}$ with $\mathbf{b}' \mathbf{1} = 1$, and $D(\mathbf{p}) = \mathbf{d}' \ln \mathbf{p}$ with $\mathbf{d}' \mathbf{1} = 0$. See, for example, Lewbel (1997).⁹

For each good j , the resulting demand model is

$$q_{ji} = Q_j(\mathbf{p}, x_i, \bar{\mathbf{q}}_g, \mathbf{z}_i) + v_{jg} + u_{ji}, \quad (8)$$

where each Q_j function is given by

$$\begin{aligned} Q_j(\mathbf{p}, x_i, \bar{\mathbf{q}}_g, \mathbf{z}_i) &= (x_i - \mathbf{p}^{1/2'} \mathbf{R} \mathbf{p}^{1/2} - \mathbf{p}' \mathbf{A} \bar{\mathbf{q}}_g - \mathbf{p}' \mathbf{C} \mathbf{z}_i)^2 e^{-\mathbf{b}' \ln \mathbf{p}} \frac{d_j}{p_j} \\ &+ (x_i - \mathbf{p}^{1/2'} \mathbf{R} \mathbf{p}^{1/2} - \mathbf{p}' \mathbf{A} \bar{\mathbf{q}}_g - \mathbf{p}' \mathbf{C} \mathbf{z}_i) \frac{b_j}{p_j} + R_{jj} \\ &+ \sum_{k \neq j} R_{jk} \sqrt{p_k / p_j} + \mathbf{A}'_j \bar{\mathbf{q}}_g + \mathbf{C}'_j \mathbf{z}_i. \end{aligned} \quad (9)$$

Here, \mathbf{A}'_j is row j of \mathbf{A} and \mathbf{C}'_j is row j of \mathbf{C} . These quantity demand functions are quadratic in the budget x_i .¹⁰

In our data, prices vary geographically by state, but are fixed within each group, so we can subscript prices by g .¹¹ More generally, our model would permit observing groups

monly assumed rank three models in empirical practice are quadratic (see the above references and the Quadratic Expenditure System of Pollak and Wales (1978)). The resulting class of indirect utility functions that yield rank three, quadratic in x demand functions are those given by equation (6).

⁹To avoid multicollinearity, in our application we restrict \mathbf{R} to be diagonal. Since $J \leq 3$, our model remains Diewert-flexible in own and cross-price effects.

¹⁰There is one straightforward extension to the demand model that we consider in some of our estimates, but do not include above to save on notation. We allow a few discrete group-level characteristics (such as religion dummies) to interact with $\bar{\mathbf{q}}_g$, thereby allowing \mathbf{A} to vary with these characteristics. Identification follows immediately from identification of the model with \mathbf{A} constant, since the the same assumptions used to identify the above model with fixed \mathbf{A} can be applied separately for each value of these characteristics.

¹¹More generally, our model would permit observing groups in multiple time periods, with prices varying by time instead of, or in addition to, varying geographically. In the Appendix (Lewbel, Norris, Pendakur, and Qu (2022)), we derive results at this added level of generality, including t subscripts for time and price regimes.

in multiple time periods, with prices varying by time instead of, or in addition to, varying geographically. In the Appendix, we derive results at this added level of generality, including t subscripts for time and price regimes.

As is standard in the estimation of continuous demand systems, we only need to estimate the model for goods $j = 1, \dots, J - 1$. The parameters for the last good J are then obtained from the adding up identity that $q_{Ji} = (x_i - \sum_{j=1}^{J-1} p_j q_{ji}) / p_J$. While we report some results using $J = 3$ goods, most of our analyses will be based on $J = 2$, with the two goods being food and nonfood. In this case, $J - 1 = 1$, so we only need to estimate the demand equation for one good, which we choose to be food. Most of our analyses will also assume \mathbf{A} and \mathbf{R} are diagonal. With these simplifications, equation (9) reduces to the single equation

$$Q_1(\mathbf{p}_g, x_i, \bar{\mathbf{q}}_g, \mathbf{z}_i) = X_i^2 e^{-(b_1 \ln p_{1g} + (1-b_1) \ln p_{2g})} d_1 / p_{1g} + X_i b_1 / p_{1g} + R_{11} + A_{11} \bar{q}_{g1} + \mathbf{C}'_1 \mathbf{z}_i,$$

where

$$\begin{aligned} X_i &= X(\mathbf{p}_g, x, \bar{\mathbf{q}}_g, \mathbf{z}_i) \\ &= x_i - R_{11} p_{1g} - R_{22} p_{2g} - (A_{11} \bar{q}_{g1} + \mathbf{C}'_1 \mathbf{z}_i) p_{1g} - (A_{22} \bar{q}_{g2} + \mathbf{C}'_2 \mathbf{z}_i) p_{2g}. \end{aligned} \quad (10)$$

As is common in empirical work in demand analysis, we recast quantity demand equations as spending equations by multiplying by price. Substituting the above into (8) and multiplying by p_{1g} yields our primary estimation model:

$$\begin{aligned} p_{1g} q_{1i} &= X_i^2 e^{-(b_1 \ln p_{1g} + (1-b_1) \ln p_{2g})} d_1 + X_i b_1 + R_{11} p_{1g} + A_{11} p_{1g} \bar{q}_{g1} \\ &\quad + \mathbf{C}'_1 p_{1g} \mathbf{z}_i + p_{1g} v_{1g} + p_{1g} u_{1i} \end{aligned} \quad (11)$$

The goal will be estimation of the set of parameters $\{\mathbf{A}, \mathbf{C}, \mathbf{R}, \mathbf{d}, \mathbf{b}\}$. In particular, \mathbf{A} embodies the impact of peer effects on needs, and hence on social welfare.

3. IDENTIFICATION AND ESTIMATION: ECONOMETRIC ISSUES

There are many obstacles to identifying and estimating our model. These issues stem from: (1) model nonlinearity (which arises from utility maximization); (2) identifying the effect of a group level variable $\bar{\mathbf{q}}_g$ in the presence of group level fixed or random effects \mathbf{v}_g ; (3) the possible absence of an equilibrium among group members; (4) endogeneity of $\bar{\mathbf{q}}_g$ (as in the Manski (1993) reflection problem); and (5) $\bar{\mathbf{q}}_g$ not being directly observed nor consistently estimated, because the data only contain a small number of members of each group. Although we solve all 5 issues, most of our econometric novelty relates to how we deal with issue 5, measurement error in the group means.

To illustrate how we overcome these econometric issues, we first consider a very simple model that suffers from all these same problems. Below we show informally how we identify and estimate this simple generic model. In the Appendix, we provide formal proofs of our identification method and associated estimator asymptotics, for both a multivariate extension of this generic model, and for our full consumer demand model given by equation (11).

Our model starts with cross-section data, where each observed individual i is assumed to be in a peer group $g \in \{1, \dots, G\}$. The number of peer groups G is large, so we assume $G \rightarrow \infty$. In our data, we will only observe a small number n_g of the individuals who are actually in each peer group g , so asymptotics assuming $n_g \rightarrow \infty$ (or assuming that n_g grows to the total number of people in each group) are inapplicable. We therefore assume n_g is fixed and does not grow with the sample size.

The generic model relates a scalar outcome y_i for person i in group g to \bar{y}_g , where $\bar{y}_g = E(y_j | j \in g)$, so \bar{y}_g is the population mean value of y_j over all people j in person i 's peer group g . For simplicity, assume there is a single scalar covariate x_i that affects y_i (we extend the generic model to vectors of y_i and x_i in the Appendix).

A typical peer specification with such data would be linear, for example, $y_i = \bar{y}_g a + x_i b + u_i$, where u_i is an error term uncorrelated with x_i , and the pair of constants (a, b) are parameters to estimate (see, e.g., Manski (1993, 2000) and Brock and Durlauf (2001)). However, to account for the nonlinearity and heterogeneity issues associated with our demand model, consider the more general specification

$$y_i = (\bar{y}_g a + x_i b)^2 d + (\bar{y}_g a + x_i b) + v_g + u_i, \quad (12)$$

where the term v_g is a group level fixed or random effect, and the constants (a, b, d) are the parameters to identify and estimate. We are not claiming that the functional form of equation (12) is in some way fundamental. Rather, it is just a simple nonlinear specification that nests the standard linear model as a special case, resembles our full demand model, and can be used to demonstrate all the issues (and solutions) associated with identification and estimation of our demand model. Equation (12) differs from the linear model both by the squared index term and by including a group-level fixed or random effect v_g .

We only have survey data with a modest number of observations for each group, so we do not assume we can observe the true \bar{y}_g even asymptotically. We therefore replace \bar{y}_g with an estimate \hat{y}_g making equation (12) equal to

$$y_i = (\hat{y}_g a + x_i b)^2 d + (\hat{y}_g a + x_i b) + v_g + u_i + \varepsilon_{gi}, \quad (13)$$

where the difference between \hat{y}_g and \bar{y}_g results in the additional error term ε_{gi} . By construction, ε_{gi} is given by

$$\varepsilon_{gi} = (\bar{y}_g^2 - \hat{y}_g^2) a^2 d + 2(\bar{y}_g - \hat{y}_g) x_i a b d + (\bar{y}_g - \hat{y}_g) a. \quad (14)$$

Inspection of equations (13) and (14) shows many of the obstacles to identifying and estimating the model parameters a , b , and d . First, with either fixed or random effects, v_g could be correlated with \hat{y}_g . Second, since n_g does not go to infinity, if \hat{y}_g contains y_i then \hat{y}_g will correlate with u_i . Third, again because n_g is fixed, ε_{gi} does not vanish asymptotically, and is by construction correlated with functions of \hat{y}_g and x_i . We can think of $(\bar{y}_g - \hat{y}_g)$ and $(\bar{y}_g^2 - \hat{y}_g^2)$ as measurement errors in \bar{y}_g and \bar{y}_g^2 , leading to the standard problem that mismeasured regressors are correlated with errors in the model.

The primary obstacle to identification and estimation is dealing with the above correlations between covariates and the unobservables v_g , u_i , and ε_{gi} . In contrast, two additional problems that are common in social interactions and network models will be more readily overcome. One is the Manski (1993) reflection problem, which does not arise here primarily because the group mean of x_i does not appear in the model.¹² Another possible problem is that the model might not have an equilibrium. For example, it could be that some members increasing their spending by one dollar would cause others to spend more by two dollars, making the original members feel the need to increase further to three dollars, etc. In the Appendix, we show that a single inequality ensures existence of an equilibrium. Roughly, an equilibrium exists as long as the peer effects are not too large.

We employ two somewhat different methods for identifying and estimating this model, depending on whether each v_g is assumed to be a fixed effect or a random effect. For each case, we construct a set of moment conditions that suffice to identify the coefficients, and are used for estimation via GMM.

3.1 Generic model estimation with group level fixed effects

In the fixed-effects model, we make no assumptions about how v_g may correlate with other covariates (including \bar{y}_g) or about how v_g might vary over time. Identification and estimation will therefore require removing these fixed effects in some way. As a result, identification will depend on the nonlinearity of demand, and so we must assume that $d \neq 0$. In contrast, our later random-effects model will make additional assumptions regarding v_g , but will be applicable to any linear or quadratic specification.

To remove the fixed effect v_g , we begin by differencing the outcomes of two consumers i and i' observed in in the same group g (and, if we have time variation, in the same time period). In addition, to remove some correlation issues, we define the leave-two-out group mean estimator

$$\hat{y}_{g,-ii'} = \frac{1}{n_g - 2} \sum_{l \in g, l \neq i, i'} y_l.$$

This $\hat{y}_{g,-ii'}$ is just the sample average of y for everyone who is observed in group g in the given time period, except for the individuals i and i' . Replacing \hat{y}_g in equation (13) with $\hat{y}_{g,-ii'}$, and differencing equation (13) between the individuals i and i' gives

$$\begin{aligned} y_i - y_{i'} &= 2\hat{y}_{g,-ii'}(x_i - x_{i'})abd + (x_i^2 - x_{i'}^2)b^2d + (x_i - x_{i'})b + u_i - u_{i'} \\ &\quad + \varepsilon_{gi} - \varepsilon_{gi'}, \end{aligned} \tag{15}$$

¹²The group mean \bar{x}_g does not appear in our model because our underlying utility theory of revealed preference with needs only gives rise to inclusion of group quantities (corresponding to \bar{y}_g in the generic model). When v_g is a fixed effect the reflection problem could still arise, in that v_g could be correlated with \bar{x}_g , but in that case we exploit the nonlinear structure of our model to overcome this issue. See the Appendix for details.

where

$$\varepsilon_{gi} - \varepsilon_{gi'} = 2(\bar{y}_g - \widehat{y}_{g,-ii'})(x_i - x_{i'})abd. \quad (16)$$

We can then show that, with some standard regression assumptions (see Theorem 1 in the Appendix) that

$$E(u_i - u_{i'} + \varepsilon_{gi} - \varepsilon_{gi'} \mid x_i, x_{i'}) = 0, \quad (17)$$

which we can then use to construct some of the moments needed for estimation of equation (15).

The intuition for this result can be seen by reexamining the obstacles to identification listed earlier. The correlation of v_g with \bar{y}_g , and hence with $\widehat{y}_{g,-ii'}$ does not matter because v_g has been differenced out. The leave-two-out average $\widehat{y}_{g,-ii'}$ does not correlate with u_i or $u_{i'}$ because individuals i and i' are omitted from the construction of $\widehat{y}_{g,-ii'}$. Finally, $\varepsilon_{gi} - \varepsilon_{gi'}$ is linear in $x_i - x_{i'}$, with a coefficient that can be shown to be conditionally mean zero.

Equation (15) contains functions of $\widehat{y}_{g,-ii'}$, x_i , and $x_{i'}$ as regressors, and equation (17) shows that we can use functions of x_i and $x_{i'}$ as instruments. However, we still require an instrument for $\widehat{y}_{g,-ii'}$, because of its correlation with $\varepsilon_{gi} - \varepsilon_{gi'}$. Since each y depends on x , an obvious candidate instrument for an average of y 's in a group (i.e., $\widehat{y}_{g,-ii'}$) would be an average of x 's in the group, that is, some estimate \widehat{x}_g of the mean group value \bar{x}_g . However, although $E(\varepsilon_{gi} - \varepsilon_{gi'} \mid x_i, x_{i'}) = 0$, the error $\varepsilon_{gi} - \varepsilon_{gi'}$ will in general be correlated with x_l for all observed individuals l in the group other than the individuals i and i' . Note that this problem is due to the assumption that n_g is fixed. If it were the case that $n_g \rightarrow \infty$, then we'd have $\varepsilon_{gi} - \varepsilon_{gi'} \rightarrow 0$, and this problem would asymptotically disappear.

To overcome this final obstacle to identification in the fixed-effects model (finding an instrument for $\widehat{y}_{g,-ii'}$), we require some other source of group level information. One possible source is repeated cross-section data, which are typically available in consumption surveys. Usually the same consumers are not sampled more than once (so no panel data is available), but we may have observations of other consumers in the same group from different time periods. It does not matter that these other consumers may or may not have the same fixed effects v_g or the same mean expenditures \bar{y}_g as in our main sample. All we need is an exogeneity assumption that each x_i is independent of the idiosyncratic error $u_{i'}$ of every person i' in person i 's group, and that the sample group averages \widehat{x}_g are autocorrelated over time (see the derivation of Theorem 1 in the Appendix for details). We can then take functions of these observations of \widehat{x}_g from other time periods to be our instruments for the corresponding functions of $\widehat{y}_{g,-ii'}$ that are in our model.

If survey data are only available for a single cross-section, another possibility is to further subdivide $\widehat{y}_{g,-ii'}$ into averages of two disjoint subgroups of group members, replacing $\widehat{y}_{g,-ii'}$ in equation (15) with one subgroup, and then using the other as an instrument.¹³ Alternatively, even if survey data is only available for a single cross-section,

¹³We thank the editor for suggesting this possibility. We do not pursue it further here because of the added complexity it entails, and because, unlike the use of \widehat{x}_g from other time periods, this alternative does

other data sets could provide the required group level instruments. For example, if x_i is a demographic variable, then instead of observing individuals from the same group in another time period, we could use census data to provide an estimate of \widehat{x}_g . Similarly, if x_i is a consumption budget as in our application, then average group level income data from wage or income surveys could suffice. It is not even necessary that we observe the exact same groups in other time periods or surveys. All we need is some overlap between the group definition in our main data and in the data used to construct the instrument, and some correlation between the variable used to construct the instrument and x .

Let \mathbf{r}_g denote a scalar or vector of the above described group level instruments. Let $\mathbf{r}_{gii'}$ denote the vector of $x_i, x_{i'}, \mathbf{r}_g$, and squares and cross products of these variables. We then obtain the unconditional moments

$$E[(y_i - y_{i'} - 2\widehat{y}_{g,-ii'}(x_i - x_{i'}))abd - (x_i^2 - x_{i'}^2)b^2d - (x_i - x_{i'})b\mathbf{r}_{gii'}] = 0. \quad (18)$$

Based on equation (18), the parameters a, b , and d can now be estimated using Hansen's (1982) GMM estimator. Each observation consists of a pair of individuals observed in a given group, so our sample becomes all such pairs i and i' . The estimator is equivalent to linearly regressing each pair $y_i - y_{i'}$ on the variables $\widehat{y}_{g,-ii'}(x_i - x_{i'})$, $(x_i^2 - x_{i'}^2)$, and $(x_i - x_{i'})$, using GMM with instruments $\mathbf{r}_{gii'}$, and then recovering the parameters a, b , and d from the estimated coefficients. By construction, the errors in this model are correlated across the pairs of individuals within each group, so we must cluster standard errors at the group level to obtain proper inference.

Theorem 1 in Appendix A.2 describes these results formally, including extending this model to allow for vector \mathbf{x}_i , providing formal conditions for proving that an equilibrium exists, and showing that the parameters of the model are identified by GMM using these moments. We then further extend this result in Appendix A.3 to allow for a J vector of outcomes \mathbf{y}_i , replacing the scalar a with a J by J matrix of own and cross-equation peer effects. Theorem 2 in Appendix A.5 then gives a final extension of these results, showing identification, consistent estimation, and inference of our full utility-derived demand model, given by equations (8) and (9) for each good j .

3.2 Generic model estimation with group level random effects

A drawback of the fixed-effects estimator is that differencing across individuals, which was needed to remove the fixed effects, results in a substantial loss of information. In this section, we add the additional assumptions that v_g is homoskedastic and independent of x_i , and develop a more efficient random-effects estimator that does not entail differencing. This random-effects estimator does not require nonlinearity for identification, and so is still consistent when $d = 0$.

To describe the random-effects estimator, it will be convenient to rewrite equation (12) as

$$y_i = \bar{y}_g^2 a^2 d + (a + 2x_i abd)\bar{y}_g + (x_i b + x_i^2 b^2 d) + v_g + u_i. \quad (19)$$

not extend to the random-effects model, both because of the presence of the squared y term, and because the instrument in this case would be invalid by correlating with the random effect.

As before, we need to replace the unobserved \bar{y}_g with some estimate, and this replacement will add an additional epsilon term to the errors. However, in the fixed-effects case, when we pairwise differenced this model, the quadratic term \bar{y}_g^2 dropped out. Now, since we are not differencing, we must cope not just with estimation error in \bar{y}_g , but also in \bar{y}_g^2 .

To obtain valid moments for identification now, we employ a variant of the method we used before. Again let i' denote an individual other than i in group g , construct $\hat{y}_{g,-i'}$ as before, and again replace \bar{y}_g with $\hat{y}_{g,-i'}$. The problem now is that the term $\hat{y}_{g,-i'}^2 - \bar{y}_g^2$ in ε_{gi} is not differenced out, and this term would in general be correlated with x_l for every individual l in the group, including i and i' .

To circumvent this problem, we replace the linear term \bar{y}_g with the estimate $\hat{y}_{g,-i'}$ as before, but we now replace the squared term \bar{y}_g^2 with $\hat{y}_{g,-i'} y_{i'}$. This latter replacement might seem problematic, since a single individual's $y_{i'}$ provides a very crude estimate of \bar{y}_g . However, we repeat this construction for every individual i' (other than i) in the group, and use the GMM estimator to combine the resulting moments over all individuals i' in g , thereby once again exploiting all of the information in the group. With this replacement, equation (19) becomes

$$y_i = \hat{y}_{g,-i'} y_{i'} a^2 d + (a + 2x_i abd) \hat{y}_{g,-i'} + (x_i b + x_i^2 b^2 d) + v_g + u_i + \tilde{\varepsilon}_{gii'},$$

where by construction the error $\tilde{\varepsilon}_{gii'}$ has the form

$$\tilde{\varepsilon}_{gii'} = (\bar{y}_g^2 - \hat{y}_{g,-i'} y_{i'}) a^2 d + (a + 2x_i abd)(\bar{y}_g - \hat{y}_{g,-i'})$$

In Appendix A.4, we show that $E(\tilde{\varepsilon}_{gii'} | x_i, \mathbf{r}_g) = -da^2 \text{Var}(v_g)$ and so equals a constant. Our constructions in estimating the group mean eliminates correlation of the error $\tilde{\varepsilon}_{gii'}$ with x_i . But $\tilde{\varepsilon}_{gii'}$ still does not have conditional mean zero, because both $\hat{y}_{g,-i'}$ and $y_{i'}$ contain v_g , so the mean of the product of $\hat{y}_{g,-i'}$ and $y_{i'}$ includes the variance of v_g .

It follows from these derivations that

$$E[y_i - \hat{y}_{g,-i'} y_{i'} a^2 d - (a + 2x_i abd) \hat{y}_{g,-i'} - (x_i b + x_i^2 b^2 d) - v_0 | x_i, \mathbf{r}_g] = 0, \quad (20)$$

where $v_0 = E(v_g) - da^2 \text{Var}(v_g)$ is a constant to be estimated along with the other parameters, and \mathbf{r}_g are the same group level instruments we defined earlier. Letting \mathbf{r}_{gi} be functions of x_i and \mathbf{r}_g (such as x_i , \mathbf{r}_g , x_i^2 , and $x_i \mathbf{r}_g$), we immediately obtain unconditional moments

$$E[(y_i - \hat{y}_{g,-i'} y_{i'} a^2 d - (a + 2x_i abd) \hat{y}_{g,-i'} - (x_i b + x_i^2 b^2 d) - v_0) \mathbf{r}_{gi}] = 0, \quad (21)$$

which we can estimate using GMM exactly as before, treating every pair of individuals in each group as observations and clustering standard errors at the group level.

The fixed-effects model is not identified if $d = 0$, that is, if the model is linear. In contrast, with random effects, if $d = 0$ then the model is still identified, and the above estimator (including estimation of d) will still be consistent. However, if we know a priori that $d = 0$, then a much simpler estimator could be used instead. If we know and impose that $d = 0$, then in the random-effects model observations can just be individuals rather

than pairs, one may simply take \widehat{y}_g to be the observed within group sample average, and use x_i and \mathbf{r}_g as instruments for identification and estimation.

As with the fixed-effects model, in the Appendix we extend the above quadratic random-effects model to allow for a vector of covariates \mathbf{x}_i , and to allow for a J vector of outcomes \mathbf{y}_i , replacing the scalar a with a J by J matrix of own and cross-equation peer effects. Appendix A.4 provides the formal proof of identification and associated GMM estimation for the random-effects generic model as discussed above (and for the extension to multiple equations), and Appendix A.6 proves that this identification and estimation extends to our full utility-derived demand model with random effects.

4. EMPIRICAL RESULTS

4.1 Data

For our main empirical analysis, we use household consumption data from the 61st round of the National Sample Survey (NSS) of India, which was conducted from July 2004 to June 2005. This survey contains information on household demographics and spending for a representative sample of the country.

To define appropriate peer groups, we exploit a property of multistage sampling, which is a standard feature of the NSS and other consumption surveys. To cut down on surveying costs, consumers are sampled from small geographic areas like villages and neighborhoods. These areas are particularly small and relevant in urban areas, where they are constructed to be compact and bounded by well-defined, clear-cut natural boundaries whenever possible, and so generally correspond to recognizable neighborhoods (NSS (2019)). Households in the same neighborhood are likely to be similar to each other in observable and unobservable ways because of assortative geographic selection, and are likely to be in at least indirect contact. This makes them appropriate candidates for defining our groups, and crucially are available as a byproduct of the sampling design in many consumption surveys.

We restrict our attention to urban households, where the geographic sampling areas are particularly small. Each subblock, the smallest geographic unit available in the data, has a population of roughly 150 to 400 households. In each subblock in our data, up to 10 households are sampled. We call this level of geography the *neighborhood*. To reflect the fact that much social activity is within religion and caste groups, we interact the neighborhoods with indicators of religion (Hindu or not) and caste (NSS scheduled caste/tribe or not). We refer to these groups defined by neighborhood, religion, and caste as *neighborhood-subcastes*, and use them as the peer groups in our analysis.¹⁴

Our sample includes all urban households in groups where we observe at least three households, the minimum required for our method of identification and estimation. To avoid expenditure outliers, we include only households that are between the 1st and 99th percentiles of household expenditure in each state. We also restrict our sample to

¹⁴The NSS contains information on whether the household is in a scheduled caste or tribe, but not the exact subcaste. However, since subcastes are typically geographically concentrated, we expect that the neighborhood-religion-scheduled caste groups will mostly capture subcastes as well.

TABLE 1. Summary statistics for consumption data.

	Observations ($N = 24,757$)				Pairs ($N = 128,640$)			
	Mean	SD	Min	Max	Mean	SD	Min	Max
x_i	0.99	0.59	0.072	4.7	1	0.59	0.072	4.7
q_i food	0.44	0.21	0	2	0.44	0.21	0	2
q_i nonfood	0.44	0.33	0.0069	2.7	0.44	0.32	0.0069	2.7
$\hat{q}_{g,ii'}$ food					0.44	0.15	0.027	1.7
$\hat{q}_{g,ii'}$ nonfood					0.44	0.24	0.02	2.4
p food	1.1	0.08	0.94	1.3	1.1	0.083	0.94	1.3
p nonfood	1.2	0.11	0.94	1.5	1.2	0.12	0.94	1.5
(Household size -1)/10	0.38	0.21	0	1.1	0.38	0.21	0	1.1
Age (household head, in 10 years)	0.39	0.11	0.17	0.82	0.4	0.11	0.17	0.82
Household head married	0.84	0.36	0	1	0.84	0.36	0	1
Log land owned	0.15	0.35	0	2.3	0.16	0.35	0	2.3
Ration card	0.14	0.35	0	1	0.13	0.34	0	1
Literate but no HS	0.46	0.5	0	1	0.47	0.5	0	1
High school or greater	0.26	0.44	0	1	0.26	0.44	0	1

Note: Table reports summary statistics for estimation sample.

households with 12 or fewer members, whose head is aged 20 or more. Together, these restrictions drop roughly 4% of the sample.

Table 1 shows summary statistics for our sample. The number of observed households in each group averages around 5 (with a range from 3 to 10), which is a small share of the several hundred households that comprise each group in the population. These small within group samples illustrate the importance of showing identification and consistent estimation without assuming that many of the members of each group are observed.

For our main sample, we have a total of 4599 distinct groups, and 24,757 distinct households. Our estimators use all unique household-pairs within each group, and we have a total of 128,640 such pairs.

The NSS collects item-level household spending and quantities for a large number of items. We consider only nondurable consumption items, and compute total expenditure x_i as the sum of spending on these goods. Our main results use a two-goods demand system of food and nonfood. On average 47% of nondurable expenditure is on food. An alternative specification we consider uses a three-goods demand system of food, fuel, and other.

For instruments, we use three other survey rounds (the 59th, 60th, and 62nd) to construct the neighborhood average expenditure in other years, $\hat{x}_{g,-t}$, where $-t$ denotes years (survey rounds) other than the one that our model is estimated using. We use functions of $\hat{x}_{g,-t}$ as instruments for neighborhood average food and nonfood consumption $\hat{y}_{g,-ii'}$. However, since the neighborhood identifiers in the NSS are not consistent over time (and are not linked to external information like the neighborhood name), we cannot identify the exact same neighborhoods in other years. For each group g , we therefore construct $\hat{x}_{g,-t}$ using all observations from other years in the same district as g . As

discussed earlier, these remain valid as instruments as long as they include some other members from the group g in other years. Our groups are spread across 535 districts, which are subunits of 20 states.

We construct prices of our demand aggregates at the state level, following Deaton (1988). We first compute state-item average unit-value prices for the subset of items for which we have quantity data. Then, in a second stage, we aggregate these state-item-level unit value prices into state-level food and nonfood prices using a Stone price index, with weights given by the overall sample average spending on each item.¹⁵

We condition on 7 demographic variables \mathbf{z} . These are household size minus 1 divided by 10; the age of the head of the household divided by 120; an indicator that there is a married couple in the household; the natural log of one plus the number of hectares of land owned by the household; an indicator that the household has a ration card for basic foods and fuels; and indicators that the highest level of education of the household head is primary or secondary level (equaling zero for uneducated or illiterate household heads).

Table 1 shows summary statistics at the level of the household, and at the level of the household-pairs used for estimation. Total expenditures and the spending components are expressed in units of average household expenditure. Only 26% of households have at least a high school education, and almost all households have married household heads. Roughly 14% of households have ration cards entitling them to subsidized basic foods.

4.2 Generic model

Our demand model assumes that the effects of peer expenditures on utility have observable implications in the corresponding demand functions (via Roy's identity). This could be violated if, for example, utilities were additively separable in $\bar{\mathbf{q}}_g$ and \mathbf{q} .

Before proceeding with our main structural results, we implement the simpler generic model of equation (15) to examine these key assumptions. Details of the data construction and empirical results of these preliminary data analyses are given in Appendix B. Here, we just briefly summarize our main findings from these empirical analyses.

We use the same data and group definitions as in our main analysis, and similarly let y_i equal expenditures on food and x_i equal total household expenditure. We report the main results of this analysis in Table A1 in the Appendix. We confirm that that peer-average food expenditures significantly affects demand for food, and that both linear and quadratic terms in the budget x_i are statistically significant. The estimated peer effects in the generic fixed-effects model are relatively imprecise, in part because the generic model does not exploit all the restrictions inherent in the structural demand model. We discuss these preliminary results in full in the Appendix, Section B.1.

¹⁵In a typical state, these prices are computed as averages of roughly 2000 observations. Given this relatively large number of observations, we do not attempt to instrument for possible remaining measurement errors in these constructed price indices.

4.3 Baseline model

Our baseline structural model is a two-goods demand system (food vs. other nondurable expenditure), as given by equation (11), and estimated by GMM using the associated moment conditions (18) and (21) for fixed and random effects, respectively. Both models use pairwise data based on all unique pairs of observations within each group, with standard errors clustered at the district level to obtain valid inference.¹⁶

Our fixed-effects approach involves substituting the leave-two-out within-group sample average quantity $\widehat{q}_{gj,-ii'}$ for the within-group mean \bar{q}_{gj} , and differencing across people within groups. Thus, we substitute $\widehat{q}_{gj,-ii'}$ for \bar{q}_{gj} in the definition of X_i (equation (10)) to create \widehat{X}_i as

$$\widehat{X}_i = x_i - R_{11}p_{1g} - R_{22}p_{2g} - (A_{11}\widehat{q}_{g1,-ii'} + \mathbf{C}'_1\mathbf{z}_i)p_{1g} - (A_{22}\widehat{q}_{g2,-ii'} + \mathbf{C}'_2\mathbf{z}_i)p_{2g},$$

and substitute $\widehat{q}_{gj,-ii'}$ for \bar{q}_{gj} and \widehat{X}_i for X_i in the demand equation (11). Then we difference the demand equation across individuals within groups to generate a moment condition analogous to (18):

$$\begin{aligned} E[(p_{1g}q_{1i} - p_{1g}q_{1i'} - (\widehat{X}_i^2 - \widehat{X}_{i'}^2)e^{-(b_1 \ln p_{1g} + (1-b_1) \ln p_{2g})} d_1 \\ - (\widehat{X}_i - \widehat{X}_{i'})b_1 - \mathbf{C}'_1 p_{1g}(\mathbf{z}_i - \mathbf{z}_{i'}))\mathbf{r}_{gi'}] = 0. \end{aligned} \quad (22)$$

Notice that, as in the generic model, many group-varying terms, including $A_{11}p_{1g}\bar{q}_{g1}$, drop out as a result of this differencing. Further, since $(\widehat{X}_i - \widehat{X}_{i'}) = x_i - x_{i'} - \mathbf{C}'_1(\mathbf{z}_i - \mathbf{z}_{i'})p_{1g} - \mathbf{C}'_2(\mathbf{z}_i - \mathbf{z}_{i'})p_{2g}$, such variables are present only in the quadratic term $(\widehat{X}_i^2 - \widehat{X}_{i'}^2)$ via interactions between group-average quantities \bar{q}_{g1} and other elements of \widehat{X}_i (e.g., x_i). The formal derivation of these moments for GMM estimation is given in Appendix A.5.

Our random-effects approach, derived in Appendix A.6, involves substituting the within-group sample average quantity and another group member's quantity for the within-group means. We use the above definition of \widehat{X}_i for the linear term in the demand equation (11) and compute a new variable $\widetilde{X}_{ii'}$ for the squared term as follows:

$$\widetilde{X}_{ii'} = \widehat{X}_i[x_i - R_{11}p_{1g} - R_{22}p_{2g} - (A_{11}q_{g1i'} + \mathbf{C}'_1\mathbf{z}_i)p_{1g} - (A_{22}q_{g2i'} + \mathbf{C}'_2\mathbf{z}_i)p_{2g}].$$

Finally, we substitute $\widehat{q}_{g1,-ii'}$ for \bar{q}_{gj} , \widehat{X}_i for X_i and $\widetilde{X}_{ii'}$ for X_i^2 in the demand equation (11) to generate a moment condition analogous to (21):

$$\begin{aligned} E[(p_{1g}q_{1i} - \widetilde{X}_{ii'}e^{-(b_1 \ln p_{1g} + (1-b_1) \ln p_{2g})} d_1 - \widehat{X}_i b_1 - R_{11}p_{1g} - A_{11}p_{1g}\widehat{q}_{g1,-ii'} \\ - \mathbf{C}'_1 p_{1g}\mathbf{z}_i - p_{1g}v_0)\mathbf{r}_{gi}] = 0. \end{aligned} \quad (23)$$

These moments use pair-specific instruments that differ between our fixed- and random-effects models. As discussed earlier, to instrument for \widehat{q}_{gj} , we construct group

¹⁶The fact that we use pairwise estimation within groups implies that we should cluster no smaller than the group level. However, because the instruments are computed at the district level, we cluster at the larger level of the district. Typical districts contain about 10 groups.

averages at the district level from other time periods. Recall that the subscript $-t$ indicate averages from all other time periods. For both the fixed and random-effects models, we create a group-level instrument \check{q}_{gj} equal to the OLS predicted value of \hat{q}_{gj} conditional on $\hat{x}_{g,-t}$, $\hat{x}_{g,-t}^2$, $\sqrt{\hat{x}_{g,-t}}$, $\hat{x}_{g,-t}^2$, $\hat{\mathbf{z}}_{g,-t}$.¹⁷

Let $\tilde{\mathbf{z}}_i$ and $\tilde{\mathbf{z}}_g$ be, respectively, the individually-varying and group-level subvectors of \mathbf{z}_i . In our baseline model, $\tilde{\mathbf{z}}_i$ includes all covariates; however, when we consider additional heterogeneity in peer effects, we will additionally include group-level covariates in $\tilde{\mathbf{z}}_g$. Letting \cdot denote elementwise multiplication, our complete instrument list for the fixed-effects model is

$$\mathbf{r}_{gii'} = (x_i^2 - x_{i'}^2), (x_i - x_{i'}) \cdot (1, \mathbf{p}_g \cdot \check{\mathbf{q}}_g, \mathbf{p}_g \cdot \tilde{\mathbf{z}}_g), \mathbf{p}_g \cdot (\tilde{\mathbf{z}}_i - \tilde{\mathbf{z}}_{i'}) \cdot (1, \mathbf{p}_g \cdot \check{\mathbf{q}}_g), x_i \mathbf{p}_g \cdot (\tilde{\mathbf{z}}_i - \tilde{\mathbf{z}}_{i'}).$$

Our instrument list for the random-effects model is

$$\mathbf{r}_{gi} = (1, \mathbf{p}_g, \mathbf{p}_g \cdot \check{\mathbf{q}}_g, \mathbf{p}_g \cdot \mathbf{z}_i), x_i \cdot (1, \mathbf{p}_g, x_{it}, \mathbf{p}_g \cdot \check{\mathbf{q}}_g, \mathbf{p}_t \cdot \mathbf{z}_{ig}), \mathbf{p}_g \cdot \mathbf{p}_g.$$

The last term provides instruments for v_0 in equation (20).

Our primary focus is on the peer effects given by elements of the matrix \mathbf{A} . We start with the simplest and most interpretable version of this structural model, where $\mathbf{A} = a\mathbf{I}_J$ is a diagonal matrix with the scalar a replicated in each element of the main diagonal. In this specification, an increase in the group-average food quantity of δ increases needs for food by $a\delta$, and an increase in the group-average nonfood quantity of δ increases needs for nonfood nondurables by the same $a\delta$. Also, having \mathbf{A} be diagonal means that group-average food quantities have no effect on needs for nonfood nondurables (and vice versa). We relax these restrictions later.

In this restricted version of the model, the welfare implications of peer effects simplify. Needs are given by $\mathbf{f}_i = \mathbf{A}\bar{\mathbf{q}}_g + \mathbf{C}\mathbf{z}_i$ and group-average expenditure is given by $\bar{x}_g = \mathbf{p}'\bar{\mathbf{q}}_g$, so when $\mathbf{A} = a\mathbf{I}_J$, the cost of needs, $\mathbf{p}'\mathbf{f}_i$, simplifies to $\mathbf{p}'\mathbf{f}_i = a\bar{x}_g + \mathbf{p}'\mathbf{C}\mathbf{z}_i$. Consequently, the scalar a equals the increase in the rupee cost of needs, $\mathbf{p}'\mathbf{f}_i$, of a one rupee increase in group-average expenditure \bar{x}_g .

4.4 Baseline estimates and alternative group sizes

Table 2 gives estimates of the scalar a . In our baseline model, groups are defined by neighborhood-subcastes, that is, a group is people who live in the same neighborhood, are of the same religion (either Hindu or not), and are of the same caste status (either scheduled caste or not). For comparison, we also consider two larger group sizes: people who live in the same neighborhood regardless of religion and caste, and people who live in the same district regardless of religion and caste.

Note that neighborhoods have populations of roughly 150 to 400 households, of which at most 10 are observed in our sample. Districts are much larger than neighborhoods, with populations of roughly 500,000 to 3,000,000 households. In our data, we

¹⁷We do this instead of using all the separate variables as instruments for \mathbf{q}_g to reduce the dimensionality of our instrument vector. This dimension reduction is needed for feasibility of our GMM estimator, because $\check{\mathbf{q}}_g$ is multiplied by the demographic controls to generate the final instrument vector.

TABLE 2. Estimated peer effects by group definition.

	RE			FE			Difference		
	District (1)	Neighborhood (2)	Neighborhood- Caste (3)	District (4)	Neighborhood (5)	Neighborhood- Caste (6)	District (7)	Neighborhood (8)	Neighborhood- Caste (9)
A (group consumption)	0.334 (0.044)	0.558 (0.036)	0.606 (0.036)	<i>Panel A: All data</i>					
N pairs	3,761,688	195,282	128,640	-0.228 (0.138)	0.088 (0.121)	0.266 (0.119)	0.562 (0.131)	0.470 (0.115)	0.341 (0.114)
N households	30,184	29,462	24,757	3,761,688	195,282	128,640	3,761,688	195,282	128,640
N groups	568	4282	4599	30,184	29,462	24,757	30,184	29,462	24,757
Average group size	53.14	6.88	5.38	568	4282	4599	568	4282	4599
				53.14	6.88	5.38	53.14	6.88	5.38
A (group consumption)	0.367 (0.045)	0.568 (0.034)	0.606 (0.036)	<i>Panel B: Consistent sample</i>					
N pairs	2,564,578	150,184	128,640	-0.284 (0.153)	0.156 (0.133)	0.266 (0.119)	0.651 (0.147)	0.412 (0.128)	0.341 (0.114)
N households	24,757	24,757	24,757	2,564,578	150,184	128,640	2,564,578	150,184	128,640
N groups	564	3941	4599	24,757	24,757	24,757	24,757	24,757	24,757
Average group size	43.90	6.28	5.38	564	3941	4599	564	3941	4599
				43.90	6.28	5.38	43.90	6.28	5.38

Note: Selected estimates for structural demand model. Controls include household size, age, marital status, land owned, ration card indicator, education, religion, and group size. Standard errors clustered at the district level.

observe 5.4 households from the average neighborhood-subcaste, while with the larger group definitions we average 6.9 and 53.1 observed households per group, respectively.

We report results for two samples. The upper half of Table 2 (Panel A) uses all the data available for each of the three group definitions, and so ends up with somewhat different samples for each. Panel B holds the sample constant across the group definitions, using only the observations from our baseline model (the smallest group definition).

Table 2 reports both random-effects (RE) and fixed-effects (FE) estimates of the scalar a , for all three group sizes. Columns (1) to (3) give RE, columns (4) to (6) give FE, and columns (7) to (9) give the difference RE minus FE.

A key implementation question is how to define our groups. If we define them at too large a level, we should expect the estimated peer effects to be biased toward zero, because our estimate of group consumption $\hat{q}_{gj,-it}$ will be mismeasured by including consumption from nonpeers. We should similarly expect the significance level of the estimates to fall if the defined groups are too large. In contrast, if we define our groups at too small a level, the estimator will likely be consistent but inefficient, because although we are grouping only households that do indeed have peer effects on each other, in each group we will be leaving out some informative peers who were placed in another group.

For both RE and FE, we find that the larger group sizes have estimates that are closer to zero and have lower t statistics than our baseline, suggesting that our baseline groups, while quite small, are the most appropriate size (the largest group size FE estimate actually flips the sign to negative, but is not statistically significant). We therefore focus our remaining analyses on the baseline neighborhood-subcaste group definition, reported in columns (3) and (6), and the difference between them in column (9).¹⁸

As expected, the RE estimates have far lower standard errors than the FE estimates, because they are based on much stronger assumptions, and do not lose information from differencing. The RE point estimate of 0.606 in column (3) also turns out to be much larger than the FE estimate of 0.266 in column (6), and we reject equality of the coefficients (column (9)).

Random effects imposes strong restrictions on unobserved heterogeneity that may not be valid, and that fixed effects do not impose, potentially biasing the RE estimates. In particular, our estimated positive difference between RE and FE estimates is consistent with group-level preferences for food consumption v_{gj} being correlated with group expenditure levels, causing upward bias in the RE peer effects estimates. This is easiest to see in a simplified version of equation (13). Suppose that the true model was linear (so $d = 0$), and we instrumented for \hat{y}_g only with other-period group consumption $\hat{x}_{g,-t}$. Then positive correlation between group expenditure and group tastes (conditional on x_i) would result in upwards bias in the estimated peer effects for normal goods like food.

¹⁸If our groups are appropriately defined, then household demands should not be strongly correlated with the average spending of individuals in other groups. We assess this by considering a placebo experiment where we randomly permute individuals into other peer groups. This is analogous to the strategy on sharp tests in network settings (Athey, Eckles, and Imbens (2020)). Using the FE estimator, the mean of the placebo distribution of a is -0.064 , and the upper 95th percentile of that distribution 0.083. This is far closer to zero than our baseline estimates even if demeaned (our FE baseline estimate is 0.266). These placebo test results are consistent with our model having appropriately defined and relevant groups.

In applications like ours where RE has much lower variance than FE (as indicated by standard errors) and is likely to be biased, to reduce mean squared errors it is common to employ shrinkage estimators. These are constructed as weighted averages of RE and FE estimates, trading off the bias of RE with the higher variance of FE (a recent example is Armstrong, Kolesár, and Plagborg-Møller (2020)). We report both the RE and FE estimates in our remaining empirical analyses, so one may implement such shrinkage if desired. However, for simplicity in our later policy discussions, we will focus on the smaller FE coefficients as conservative estimates of the magnitude of peer effects. Below we consider a number of robustness checks and alternative specifications. Most yield larger (but less significant) effects than our baseline FE estimate of $a = 0.266$, which we therefore take to be a conservative estimate of the magnitude of peer effects.

To interpret our estimate of a , imagine first that just one household in a group was given an additional s rupees to spend. Compare this to the case where everyone in the group each had an additional 1000 rupees to spend. What s would give the household in the first case the same utility as in the second case? The answer must be less than a 1000 rupees, because in the second case, peer effects reduce the utility of the increased spending. By our model, the answer is $s = 1000(1 - a)$, which is 734 rupees in the FE model.

Economic theory requires that a lie between zero, and roughly, one. It is greater than zero because our model is one of the peer effects increasing perceived needs that take the form of costs, and it is less than about one to ensure that an equilibrium exists.¹⁹ An encouraging feature of our estimates (both RE and FE) is that they lie well within this required range, without any such constraint being enforced in estimation.

4.5 Measurement error in group means

The neighborhood-subcaste groups in our baseline analysis each have between 3 and 10 observed households, out of an average of around 200 households in the population. This suggests that the group mean measurement errors $\hat{q}_{gj} - \bar{q}_{gj}$ are likely to be substantial. Much of the complexity in our GMM estimator entails constructing moments that remain valid in the presence of these measurement errors. Table 3 considers the impact of our measurement error corrections on the estimated values for a in both the RE and FE models. We should expect that, the smaller are the group definitions, the larger are the measurement errors in the estimates of each \bar{q}_{gj} , and hence the larger should be the effect of correcting for these measurement errors.

Regarding the direction of bias, one might expect measurement error in \bar{q} to induce the usual attenuation (i.e., bias toward zero) that is standard in linear models with measurement error. However, the nonlinearity of our models and our estimators could cause bias in either direction. A priori, we expect standard attenuation bias to play a larger role in the RE model, because in that model the parameter a is primarily identified as the coefficient of the estimate of \bar{q} itself (as in linear models), while in the FE model, due to

¹⁹The exact value that is necessary to ensure that an equilibrium exists has a complicated expression, which we derive in the Appendix, but this value is near one.

TABLE 3. Estimated peer effects by measurement error correction.

	RE			FE		
	District (1)	Neighborhood (2)	Neighborhood- Caste (3)	District (4)	Neighborhood (5)	Neighborhood- Caste (6)
<i>Panel A: Naive (no correction)</i>						
A (group consumption)	0.143 (0.031)	0.038 (0.017)	0.054 (0.016)	0.470 (0.215)	0.559 (0.089)	0.529 (0.090)
N pairs	2,564,578	150,184	128,640	2,564,578	150,184	128,640
N households	24,757	24,757	24,757	24,757	24,757	24,757
N peer groups	564	3941	4599	564	3941	4599
Average group size	43.90	6.28	5.38	43.90	6.28	5.38
<i>Panel B: Baseline</i>						
A (group consumption)	0.367 (0.045)	0.568 (0.034)	0.606 (0.036)	-0.284 (0.153)	0.156 (0.133)	0.266 (0.119)
Hausman H	49.15	314.55	295.21		16.71	11.48
<i>p</i> -value	0.000	0.000	0.000		0.000	0.001
N pairs	2,564,578	150,184	128,640	2,564,578	150,184	128,640
N households	24,757	24,757	24,757	24,757	24,757	24,757
N peer groups	564	3941	4599	564	3941	4599
Average group size	43.90	6.28	5.38	43.90	6.28	5.38

Note: Selected estimates for structural demand model. Controls include household size, age, marital status, land owned, ration card indicator, education, religion, and group size. Standard errors clustered at the district level.

differencing, a is identified only off of differences of interactions between \bar{q} and other covariates.

To assess the impact of our corrections for measurement error, we replace the instruments in our models with stronger instruments that would be valid in the absence of measurement error. In particular, instead of instrumenting \hat{q}_{gj} with district-level averages from other time periods, we instrument \hat{q}_{gj} with group-level averages from the current time period. So, everywhere that $\hat{x}_{g,-t}$ and $\hat{z}_{g,-t}$ appear in our estimators, we replace them with \hat{x}_g and \hat{z}_g . As a result, the total number and types of moments remains exactly the same as in our baseline estimates.

Table 3 is analogous to columns (1) to (6) of Table 2, but is estimated with the instruments that do not correct for measurement error. This should be compared to the corresponding entries in Table 2. Columns (3) and (6) are still our preferred group size specifications.

Both the RE and FE estimates show considerable differences between estimates with and without the measurement error correction. As expected, the smaller the group sizes, the larger the differences between the corrected and uncorrected estimates.

In the RE models, we see standard attenuation bias dominating, and the magnitude of the bias appears very large: uncorrected estimates are about half the size of the corrected estimates for the largest group size, attenuating all the way to about one-tenth the size of the corrected estimate for our baseline, which is the smallest group size. The measurement error corrected RE estimates also have larger standard errors than the uncor-

rected estimates, due to the fact that the instruments are less informative in the former case. Since both estimators would be consistent in the absence of measurement error, we can form a Hausman test to compare the estimators, and the uncorrected estimators are rejected.

The direction and size of bias is different for the FE estimator. Here, at all three group sizes, the uncorrected estimates are about twice as large as the corrected, suggesting a significant impact of nonlinearity and differencing on the size and direction of bias in the FE models. As with the RE models, the uncorrected FE estimates have smaller standard errors than the corrected estimates, and Hausman tests reject the uncorrected estimates. We conclude that our corrections for measurement errors due to small within group sample sizes are empirically justified and important.

4.6 *Alternative specifications and robustness checks*

4.6.1 *Peer effects by demographic groups* In Tables 2 and 3, the peer effect parameter a is restricted to be the same for all types of households. In Table 4, we allow a to vary with observed household characteristics. In columns (1) and (5), we replicate columns (3) and (6) from Table 2, where the group is defined at the neighborhood-subcaste level, for the RE and FE models, and a is a fixed value. In columns (2) and (6), we allow a to depend on whether the household is Hindu or not, and whether they come from a scheduled (disadvantaged) caste. In columns (3) and (7), we define groups at the neighborhood-subcaste-landownership level, and allow a to depend on the landownership indicator variable. In columns (4) and (8), we define groups by neighborhood-subcaste-high-school attainment, and allow a to depend on the high-school attainment indicator.²⁰

Columns (2) and (6) show estimated differences in peer effects across Hindu versus non-Hindu and scheduled versus nonscheduled tribe/caste. The left-out category (picked up by the constant) is nonscheduled Hindu. The RE estimates show some significant differences in peer effects, but the FE estimates do not, and most of estimated differences have the opposite sign in the FE versus RE models.

Columns (3) and (7) allow a to depend on the household level land-ownership indicator. Both the FE and RE models show landowners having larger peer effects than landless households, but the magnitudes differ dramatically, with RE estimates implying a small difference, while FE showing the landless having almost no peer effects. As before, the standard errors on the FE models are all much larger than the RE standard errors.

Columns (4) and (8) allow a to depend on a household level high-school attainment indicator, defined to equal to 1 if the household head has at least high-school education, and zero otherwise. Here, the FE and RE models disagree with the FE model showing the more educated households having larger peer effects, while the RE model shows the opposite.

Particularly, when focusing on the FE estimates, our estimated peer effects are larger for higher socioeconomic status groups. A possible explanation is that the poorest

²⁰As the groups become smaller, the number of groups actually declines, because we can include only groups with at least 3 members.

TABLE 4. Peer effects by demographic group.

	RE				FE			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Constant	0.606 (0.036)	0.606 (0.035)	0.600 (0.057)	0.694 (0.041)	0.266 (0.119)	0.255 (0.121)	0.078 (0.125)	0.145 (0.143)
Scheduled non-Hindu		0.168 (0.066)				0.130 (0.237)		
Scheduled Hindu		0.247 (0.087)				-0.285 (0.330)		
Nonscheduled non-Hindu		0.179 (0.055)				-0.075 (0.153)		
Owns land			0.031 (0.054)				0.446 (0.136)	
High school or greater				-0.186 (0.057)				0.454 (0.163)
<i>p</i> -value heterogeneity		0.00	0.56	0.00		0.68	0.00	0.01
N pairs	128,640	128,640	100,756	84,052	128,640	128,640	100,756	84,052
N households	24,757	24,757	21,696	20,233	24,757	24,757	21,696	20,233
N peer groups	4599	4599	4452	4446	4599	4599	4452	4446

Note: Selected estimates for structural demand model, controls include household size, age, marital status, land owned, ration card indicator, education, religion, and group size. A (group consumption) represents the structural effect of group consumption on own consumption, constrained to be the same for all goods. Standard errors clustered at the district level.

households in India are close enough to subsistence that it is more costly to engage in status competitions. This is similar to [Akay and Martinsson's \(2011\)](#) finding for very poor Ethiopians.

4.6.2 Cross-group peer effects Our baseline estimates allow only for within-group consumption peer effects. However, conceptually, it is possible that needs could depend on consumption levels of other “nearby” peer groups. Our baseline grouping structure is neighborhood-subcaste, so that in a given neighborhood, there could be several groups defined by varying religion and caste. In this subsection, we consider the possibility that peer effects may be relevant between groups, and that, in particular, needs may be “upward-looking” or aspirational, in the sense that perceived needs are affected by the consumption behavior of our betters in the social hierarchy. We operationalize this by focusing on a subset of 564 groups that are low-caste Hindu, and allowing for both within-group peer effects and for peer effects, which depend on the consumption of upper-caste Hindus in the same neighborhood. In this model, the cost of needs of lower-caste households is $\mathbf{p}'\mathbf{f}_i = a_l\bar{x}_{l,g} + a_u\bar{x}_{u,g} + \mathbf{p}'\mathbf{Cz}_i$, where a_l gives the effect of own-group (lower-caste) spending in the neighborhood, $\bar{x}_{l,g}$, and a_u gives the effect of upper-caste spending in the same neighborhood, $\bar{x}_{u,g}$. We present RE and FE estimates of this model in Table 5.

As with the baseline estimates, the RE estimates of within-group peer effects are much larger than the FE estimates. The RE estimate of the within-group peer effect, a_l , for these lower-caste households is 0.802, which is similar to the baseline RE estimate. The FE estimate of a_l is 0.445, which is somewhat higher than the baseline estimate of

TABLE 5. Peer effects of own-group and out-group spending.

	RE		FE	
	(1) Own-Group	(2) Upper Caste	(3) Own-Group	(4) Upper Caste
A (group consumption)	0.802 (0.070)	0.052 (0.039)	0.445 (0.138)	0.011 (0.302)
Number of pairs	8962	8962	8962	8962
Number of groups	564	564	564	564

Note: Dependent variable is household food spending. Individual controls include household size, age, marital status, and amount of land owned. All models include price controls. Analysis is restricted to scheduled castes and tribes in FSUs with at least 3 nonscheduled caste or tribe households. Own-group columns display the peer effects of own-group expenditure; upper caste columns display the peer effect of the neighbor upper caste household expenditure. Standard errors in parentheses and clustered at the group level.

0.266, but given the estimated standard error of 0.138, it is close in terms of its sampling distribution. (As with the baseline estimates, the RE model is rejected against the FE alternative.) Turning to the cross-group (within neighborhood) peer effects, these are small and statistically insignificant for both the RE and FE models. The estimated value of the cross-group effect is 0.052 and 0.011 in RE and FE models, respectively. Given that these cross-group effects are relatively small in magnitude and statistically insignificant, we conclude that our model with just within-group and not cross-group peer effects appears to suffice.

4.6.3 Peer effects with alternative specifications of the A matrix Next, Table 6 considers what happens when we relax the restriction that $\mathbf{A} = a\mathbf{I}_J$ for a scalar a . Since needs are given by $\mathbf{f}_i = \mathbf{A}\bar{\mathbf{q}}_g + \mathbf{C}\mathbf{z}_i$, the money cost of the part of needs driven by peer effects is given by $\mathbf{p}'\mathbf{A}\bar{\mathbf{q}}_g$. In the previous subsections, with $\mathbf{A} = a\mathbf{I}_J$, this cost of needs due to peer effects is $\mathbf{p}'\mathbf{A}\bar{\mathbf{q}}_g = a(p_1\bar{q}_{1g} + p_2\bar{q}_{2g}) = a\bar{x}_g$, and so is proportional to group mean total expenditures \bar{x}_g . When we allow \mathbf{A} to be an unconstrained diagonal matrix, this cost of needs becomes $\mathbf{p}'\mathbf{A}\bar{\mathbf{q}}_g = a_{11}p_1\bar{q}_{1g} + a_{22}p_2\bar{q}_{2g}$. This allows for the possibility that group-average food expenditure, $p_1\bar{q}_{1g}$, and group-average nonfood nondurable expenditure, $p_2\bar{q}_{2g}$, have different effects on needs. Finally, when \mathbf{A} is completely unrestricted, we get $\mathbf{p}'\mathbf{A}\bar{\mathbf{q}}_g = a_{11}p_1\bar{q}_{1g} + a_{21}p_2\bar{q}_{1g} + a_{12}p_1\bar{q}_{2g} + a_{22}p_2\bar{q}_{2g}$.

In columns (3) and (6) of Table 2, we reproduce columns (3) and (6) of Table 6, reporting the estimate of the scalar a where $\mathbf{A} = a\mathbf{I}_J$. In Columns (2) and (5) of Table 6, we let \mathbf{A} be an unconstrained diagonal matrix, and report its two estimated diagonal elements, a_{11} and a_{22} . And in columns (1) and (4) of Table 6, we give estimates of all four elements of \mathbf{A} where \mathbf{A} is completely unrestricted. For these estimates, we again define groups as neighborhood-subcaste.

The main difficulty in estimating these more general models is multicollinearity. As people's income rises, they tend to spend more on both food and nonfood items. As a result, $p_1\bar{q}_{1g}$ and $p_2\bar{q}_{2g}$ (group average food and nonfood expenditures, resp.) in the diagonal \mathbf{A} model tend to be highly correlated across groups. This problem is worse still in the unrestricted \mathbf{A} model, where $p_1\bar{q}_{1g}$, $p_2\bar{q}_{1g}$, $p_1\bar{q}_{2g}$, and $p_2\bar{q}_{2g}$ are all highly multicollinear, both because group-average quantities of food and nonfood are positively

TABLE 6. Peer effects by A matrix specification.

	RE			FE		
	(1)	(2)	(3)	(4)	(5)	(6)
A (group food on food consumption)	0.411 (0.171)	0.639 (0.036)	0.606 (0.036)	9.741 (2.066)	2.228 (0.382)	0.266 (0.119)
A (group nonfood on nonfood consumption)	0.452 (0.171)	0.572 (0.034)	0.606 (0.036)	5.400 (1.577)	-0.911 (0.276)	0.266 (0.119)
A (group food on own nonfood consumption)	-0.397 (0.275)			-7.695 (1.828)		
A (group nonfood on own food consumption)	-0.095 (0.102)			-6.383 (1.860)		
<i>p</i> -value equality	0.896	0.001		0.000	0.000	
<i>p</i> -value diagonal	0.002			0.000		
N pairs	128,640	128,640	128,640	128,640	128,640	128,640
N households	24,757	24,757	24,757	24,757	24,757	24,757
N peer groups	4599	4599	4599	4599	4599	4599

Note: Selected estimates for structural demand model, controls include household size, age, marital status, land owned, ration card indicator, education, religion, and group size. Standard errors clustered at the district level.

correlated with each other, and because prices are positively correlated with each other across states.

Considering first the RE estimates with an unrestricted diagonal \mathbf{A} matrix (column (2) of Table 6), we see estimated values of 0.639 and 0.572 for a_{11} and a_{22} , respectively. These are similar in magnitude to each other, and similar to the estimated value of 0.606 for a in the baseline RE model. Although the two values are similar in magnitude, they are estimated precisely enough to reject the hypothesis that they are identical. Turning to the RE estimates with an unrestricted \mathbf{A} matrix, column (1), we again find the estimated magnitudes of a_{11} and a_{22} are similar to each other (though lower than before), and the difference between them is now statistically insignificant. The off-diagonal elements of this unrestricted \mathbf{A} matrix are both statistically insignificant.

Taken together, we interpret these results as evidence that imposing the restrictions $a_{11} = a_{22}$ and $a_{12} = a_{21} = 0$, as in our baseline model, is at least a reasonable approximation.

In contrast to the RE model, we see evidence that the above discussed multicollinearity overwhelms the FE model. Column (5) shows infeasibly large estimates of a_{11} and a_{22} with opposite signs and greatly increased standard errors, and even more extreme estimates in column (4) where all four elements of \mathbf{A} have impossibly large magnitudes and varying signs. These are all common hallmarks of substantial positive multicollinearity.

We should expect that the multicollinearity issues among the $p_j \bar{q}_{kg}$ terms would be much more severe in the FE model, and not just because it is based on a weaker set of assumptions. The identification of \mathbf{A} in the FE estimator comes only from interaction terms between each $p_j \bar{q}_{kg}$ and the budget x_i . This is due to the fact that the level terms for each $p_j \bar{q}_{kg}$ get differenced away. In contrast, the identifying variation for \mathbf{A} in the RE estimator comes from both the level terms $p_j \bar{q}_{kg}$ and their interactions with x_i .

TABLE 7. Estimated peer effects in a three-goods demand system.

	RE		FE	
	(1)	(2)	(3)	(4)
A (group food on food consumption)	0.848 (0.023)	0.932 (0.014)	2.393 (0.426)	0.296 (0.100)
A (group fuel on own fuel consumption)	0.938 (0.018)	0.932 (0.014)	2.820 (0.913)	0.296 (0.100)
A (group other on own other consumption)	0.740 (0.023)	0.932 (0.014)	-1.387 (0.334)	0.296 (0.100)
Hausman H			42.151	41.701
<i>p</i> -value			0.00	0.00
<i>p</i> -value equality	0.000		0.000	
N pairs	128,640	128,640	128,640	128,640
N households	24,757	24,757	24,757	24,757
N peer groups	4599	4599	4599	4599

Note: Selected estimates for structural demand model, controls include household size, age, marital status, land owned, ration card indicator, education, religion, and group size. Standard errors clustered at the district level.

We take from these results that the multicollinearity of group-average expenditures is too severe in our data to get trustworthy estimates of variation in the elements of \mathbf{A} in our preferred fixed-effect specification; however, our baseline restriction $\mathbf{A} = a\mathbf{I}_J$ appears to be reasonable and adequate.

4.6.4 A three-goods model All of the models presented so far have been demand systems with $J = 2$ goods (food and nonfood). When $J = 2$, we only need to estimate a single demand equation (since the other is determined by the restriction that consumers exhaust their budget). However, our theorems show identification of peer effect parameters in demand systems where J is any number of goods. In Table 7, we present estimates of a $J = 3$ equation demand model, having two equations we need to estimate. The three goods are taken to be food, fuel, and other nondurable goods. The former nonfood category is now divided into fuel and other, so total expenditures x_i for each household remains the same as before.

We report estimates for the RE and FE models, with an unrestricted diagonal \mathbf{A} matrix in columns (1) and (3) of Table 7, and with the restriction that $\mathbf{A} = a\mathbf{I}_J$ in columns (2) and (4). As before, groups are defined at the neighborhood-subcaste level.

In the RE models, a in column (2) and the varying diagonal elements of \mathbf{A} in column (1) are all significant and larger than before, ranging from 0.740 to 0.938. Since adding more goods should not increase the magnitude of the overall peer effects, we take this as additional evidence that the restrictions imposed by the RE model may not hold, and are likely inducing an upward bias. We also perform a Hausman test of the RE model against the FE model, and again reject the additional restrictions imposed by the RE model.

In the FE model, we again see evidence of multicollinearity in column (3), with two elements of the estimated \mathbf{A} diagonal being extremely large and positive, and one being extremely large and negative. However, in column (4), we obtain a statistically significant estimate of a of 0.296, which agrees very well with the FE estimate of 0.266 we had in the

TABLE 8. Peer effects in spending, by consumption categorizations.

	RE			FE		
	(1) Lux	(2) Visible	(3) Vis. lux	(4) Lux	(5) Visible	(6) Vis. lux
A (group consumption)	0.545 (0.079)	0.400 (0.086)	0.654 (0.111)	0.748 (0.203)	0.418 (0.115)	0.656 (0.131)
Number of pairs	128,974	128,974	128,974	128,974	128,974	128,974
Number of groups	4607	4607	4607	4607	4607	4607

Note: Dependent variable is varies by column. Individual controls include household size, age, marital status, and amount of land owned. All models include price controls. Lux columns divide expenditure into luxuries and necessities; visible columns into visible and invisible consumption; and vis. lux into visible luxury versus other. Standard errors in parentheses and clustered at the group level.

two-goods baseline model. We take this as additional evidence in favor of the FE model with $\mathbf{A} = a\mathbf{I}_J$.

4.6.5 Alternative classifications of goods Previous research on peer effects in consumption has emphasized the possibility that such externalities may be more relevant for types of consumption that are: (a) signals of status or wealth (e.g., Veblen (1899)); or (b) easily observed (e.g., Charles, Hurst and Roussanov (2009), Heffetz (2011), and Roth (2014)). Our baseline results consider a two-goods model with demands for food and nonfood consumption. In this subsection, we consider three different alternative two-goods models, based on the whether or not the expenditures observed in the NSS are necessities or luxuries and whether they are are more versus less visible components of consumption. For the former distinction, we classify roughly 100 fine-grained consumption categories reported in the NSS as luxuries if their budget elasticity exceeds one and as necessities if less than or equal to one. In this classification, food is split, with food at home classified as a necessity and food-out classified as a luxury. For the latter distinction (visible consumption or not), we use the classification of Roth (2014, Table 4). Both food at home and food-out are classified as visible.

In Table 8, we present RE and FE estimates for three two-goods models: luxuries versus necessities (columns 1 and 4); visible versus invisible (columns 2 and 5); and visible luxuries versus not visible luxuries. We show estimates of a in a model where $\mathbf{A} = a\mathbf{I}_2$. In this table, groups are defined as our baseline neighborhood-caste level (as in columns 3 and 6 in Table 2).²¹ We do not present estimates of models with different peer effects for different goods (e.g., where \mathbf{A} is diagonal with different elements on the main diagonal), because as before the FE estimates of such models are very imprecise.

Looking first at the RE estimates, we see that the main difference between these and our baseline results is that the estimates in Table 8 are dramatically less precise. Although the point estimates are in the ballpark of the baseline estimate of 0.606, the

²¹A (much) earlier version of this paper considered these classifications of goods. But those estimates used a different, and much larger, definition of the group (district*education). Consequently, those point estimates were quite different from the the ones presented here.

estimated standard errors are roughly three times as large. That the peer effects are estimated with reduced precision suggests that these classifications of goods may yield demands that are noisier than our baseline food versus nonfood classification.

Turning to the FE estimates, we see that the point estimates are larger than in the baseline specification, and they are much closer to the RE point estimates. In fact, the Hausman test no longer rejects these RE specifications.

Note that “visible” is dominated by food (because food-at-home and food-out together make up the single-largest expenditure component, and all food expenditures are classified as visible expenditures), whereas “luxury” and “visible luxury” are dominated by food-out. For the FE models, the point estimates for “visible” are (not surprisingly) similar to our baseline based on food. Models based on “luxury” or “visible luxury” give point estimates of a that are larger than those for “visible.” The most precisely estimated of these FE models is that which contrasts visible to invisible expenditures. Here, the point estimate of a is 0.418, with an estimated standard error of 0.115. This is roughly one standard error above our baseline estimate of 0.266.

We draw three conclusions from these alternative specifications of the classification of goods. First, the demand system we choose to estimate does make a difference when it comes to the magnitude of the estimated peer effect. Second, even with these quite different classifications of goods, we find large and statistically significant peer effects for all of them. Third, given the large estimated standard errors for FE models, the general picture we obtain is similar between the baseline specification and these alternatives. Overall, our baseline FE model appears to give a conservative significant estimate of peer effects at $a = 0.266$.

4.7 *Are peer expenditures really negative externalities?*

Our findings suggest that higher peer expenditures makes consumers behave, at the margin, as if they were poorer. We take this to mean that, in a welfare sense, they feel poorer. While peer expenditures may in theory have both positive and negative effects, our model estimates therefore imply they are on average negative externalities. However, in theory, peer consumption could have alternatively increased rather than decreased the utility of the goods one consumes. An example could be something like a phone, which becomes more valuable (via network effects) when other consumers also have phones.

We address this concern by directly estimating the effect of peer expenditure on an observed measure of subjective well-being data, and confirm that, conditional on household income, higher levels of peer group expenditures are associated with lower satisfaction on average. We take this as confirmation that our demand estimates do indeed reflect lower welfare resulting from increasing peer expenditures. The full analysis is in Appendix B.2, but we provide a brief summary below.

For this exercise, we use data from the Indian modules of the World Values Surveys (WVS).²² The WVS asks respondents about their subjective well-being, and codes the

²²We use the 4th (2001), 5th (2006), and 6th (2014) waves.

response on a five-point likert scale. The WVS also includes information on household income quintiles. Since the same granular geographic identifiers are not available in the WVS, we define groups using the intersection of state and religion, and identify average expenditure for each group using the NSS data.

Interpreting ordinal self-reported well-being as a crude measure of utility, we regress this self-reported well-being on one's own income bin and on the average expenditure in one's group. The results are reported in Table A2 in the Appendix. We find that the resulting coefficient estimates have signs that are consistent with our theory: higher income increases self-reported well-being, but higher group expenditure decreases it. A 1000 rupee increase in peer group expenditure (relative to a mean of 5554, with standard deviation of 2580) decreases self-reported well-being by 15% of a standard deviation, which is in line with the welfare effects we found using our structural model.²³ As we discuss in Appendix B.2, these effects of peer expenditure are similar throughout the distribution of one's own income, consistent with our linear index structure for peer effects.

5. IMPLICATIONS FOR TAX AND TRANSFERS POLICY

Our finding that perceived needs rise with peer group average consumption has significant implications for policies regarding redistribution, transfer systems, public goods provision, and economic growth. In this section, we provide some crude, back of the envelope calculations that illustrate the rough magnitudes that our estimated peer effects have on policy questions.

Our model is one where consumption has negative externalities on one's peers. Boskin and Shoshenski (1978) consider optimal redistribution policies in models with general consumption externalities. They show that distortions due to negative externalities from consumption onto utility can generally be corrected by optimal taxation. In particular, their results imply that negative consumption externalities make the marginal cost of public funds lower than it would be otherwise. Here, we apply the same logic to our estimated consumption peer effects, and in particular, show how large free lunch gains may be possible.

A potentially peculiar attribute of our model is that it could be social welfare improving to transfer income from someone with poor peers to someone else of equal income who has rich peers. This is not a specific feature of our model; similar implications can arise as long as peer spending negatively affects individual utility. As a practical matter, we rule out such transfers, by only considering tax and transfer programs that are based

²³In principle, one could use self-reported well-being data to estimate a , the effect of peer expenditure in money-metric terms. There are three issues with this approach. First, self-reported well-being is generally crudely measured and may not be interpersonally comparable. Second, few if any existing data sets record both consumption and self-reported well-being. Third, this approach (as well as that of other papers in the literature, such as Luttmer (2005), that apply this approach) relies on a random-effects assumption that expenditures are uncorrelated with other determinants of self-reported well-being. A key advantage of our utility-derived demand model is that the FE approach allows identification even when group preferences are correlated with group expenditures. Given these issues, we take the self-reported well-being results here only as evidence of negative consumption externalities, and do not attempt to use them to back out other measures of the welfare cost of peer effects.

on personal income rather than peer group membership. Many of our conclusions then follow from the observation that the demographics that determine peer group membership (e.g., education and neighborhoods) strongly correlate with income. So, for example, transfers from high to low income households will on average transfer resources from higher socioeconomic status groups to lower status groups.

As discussed in Section 2, the sum (over households) of income minus the sum of spending on needs (as we define them) is a valid money-metric social welfare index. This means that if needs go down, all else equal, social welfare goes up. Consider the money metric costs in lost utility of, say, an across-the-board tax increase. This tax increase lowers average expenditures by households, which in turn lowers perceived needs, thereby offsetting some of the utility that was lost by having to pay the tax.

For simplicity, round our conservative baseline estimate of $a = 0.266$ to $1/4$. Suppose you experience a 4 rupee tax increase, and for simplicity let your marginal propensity to consume be 100%. If your peers also have their taxes increase by the same amount, then your loss in utility will only be equivalent to that of a 3 rupee tax increase. The reason is that although your net income and, therefore, expenditure will have dropped by 4 rupees, so will have that of your peers. Consequently, your needs will have dropped by $1/4 * 4 = 1$ rupee, so that your net loss in money-metric utility is only 3 rupees.

However, to fully evaluate the effect of this tax increase, we must also consider potential peer effects in how the government uses the additional tax revenue. If the money is transferred to other groups of consumers who also have peer effect spillovers of $a = 1/4$, then the welfare gains from reduced expenditures on needs by the taxed consumers will be offset by the welfare losses associated with increased perceived needs by the recipients of those transfers.

There are two ways we can reduce or eliminate these offsetting welfare losses, thereby exploiting the potential free lunch associated with the reduced perceived needs from taxing peers. One way is to transfer the tax revenues to individuals in groups that have smaller peer effects, and the other could be to spend the tax revenue on public goods or government services.

We found some evidence that the size of the peer effects may be smaller for poorer and less educated groups than for other consumers. If so, then transfers from higher income to lower income individuals will lead to an overall increase in social welfare, by reducing the total negative consumption externalities of peer effects. This is true even with an inequality-neutral social welfare function. Similarly, our estimates suggest social welfare gains to progressive versus flat taxes, even if the marginal utility of money was the same for all consumers.

An alternative way to exploit the potential free lunch associated with reduced perceived needs from taxing peers is to spend the resulting tax revenues on public goods or government services. To the extent that jealousy or envy are the underlying cause of the peer externalities we identify, public goods and services may not invoke those effects (or at least induce smaller peer effects), because by definition public goods are consumed by all members of the group. This suggests that public goods and services may provide at least a partially free lunch.

To illustrate the magnitude of these potential welfare gains, we consider just one existing transfer program in India. This is the Public Distribution System (PDS), which is estimated to cost roughly 1.35% of GDP when fully implemented (Puri (2017); Ministry of Consumer Affairs (2018)). The PDS aims to provide subsidized cereals to roughly 75% of Indian households. Our estimates imply that the resulting increased consumption would result in increased perceived needs, and so would not raise utility as much as an alternative policy that did not induce these negative externalities. Such alternatives could be the provision of public goods or services that provide utility to the poor but are equally available to all households. Such public goods and services might include clean water, public sanitation, better air quality, or improved fire or police protection.

A back-of-the-envelope calculation of the magnitude of these potential gains proceeds as follows. The entitlement of rice under the PDS is up to 5 kg per month per person at 3 rupees per kg. Suppose the market price of rice is 15 rupees per kg (as it was in 2016). Then the public cost of providing 5 kg of rice at the subsidized price of 3 rupees per kg is 60 rupees per month per person. Ignoring waste, the private consumption of recipient therefore increases by 60 rupees per month per person. Using $a = 1/4$, this implies that needs rise by 15 rupees per month per person. Thus, the government's expenditures of 60 rupees only increases money metric utility by 45 rupees per person per month. This is in contrast to a benefit of up to 60 rupees per person per month that might be obtained by provision of public goods. The PDS program targets roughly 1 billion people, yielding potential money-metric welfare gains (of switching from rice subsidies to a public goods program) of up to roughly 180 billion rupees (over 2 billion US dollars) per year.

This crude, toy calculation comes with many caveats,²⁴ and is therefore not a rigorous analysis of alternatives to the PDS. It is only intended to illustrate the potentially enormous impacts that accounting for peer effects could have on the evaluation of tax and transfer policies.

6. CONCLUSIONS

We show identification and consistent GMM estimation of peer effects in a model where most members of each peer group are not observed. The model allows for peer group level fixed or random effects, and allows the number of observed individuals in each peer group to be small and fixed asymptotically. This means we obtain consistent estimates of the model even though peer group means cannot be consistently estimated. Unlike most peer effects models, our model can be estimated from standard cross-section survey data where the vast majority of members of each peer group are not observed; each member is only observed once, and detailed network structure is not

²⁴Some caveats are that the benefits of this alternative might be reduced to the extent that some households derive less utility from the public good than others, or if the public good also generates negative peer effects, or if households in total derive less utility from the public good than it costs to provide. The effects might also be reduced to the extent that households given ration cards, who are therefore relatively poor, may have smaller than average peer effects, and some are likely to be at subsistence levels where few such effects are possible. On the other hand, benefits could also be increased to the extent that people in groups that did not qualify for or take up the rice entitlement might benefit from provision of the public good.

available. We obtain these results both for a generic quadratic model, and for a utility-derived demand model. The methods we use to identify and estimate these effects could potentially have broad application to other social network models.

Our estimator is designed to estimate peer effects from survey data in the absence of network information. However, components of our methodology could be useful even when network data is observed. In Appendix A.7, we show how our estimator could be applied to data where, instead of groups, each person has their own set of friends, a small subset of whom are observed. Further extensions of this type would be a useful area for future research. Another possibility for future work, especially with larger group sizes, would be to consider the possibility that peer effects are functions of statistics other than the mean (such as the variance and/or quantiles) of the within-group quantities.

We propose a utility-derived consumer demand model where a consumer's perceived needs for each good depends in part on the average consumption of goods among the other members of the consumer's peer group. We show how this model can be used for welfare analyses, and in particular to identify what fraction of total expenditure increases are spent on "keeping up with the Joneses" type peer effects.

We apply the model to consumption data from India, and find large peer effects. Our estimates imply that an increase in group-average spending of 100 rupees would induce an increase in needs of about 26 rupees or more in most peer groups. This means that the increase in utility you experience if you and everyone else in your peer group spends 100 more rupees (say, because of a tax cut) is the same as the increase in utility you would get from spending only $100 - 26 = 74$ more rupees if no one else in your peer group increased their spending.

These results can at least partly explain the [Easterlin \(1974\)](#) paradox, in that income growth over time, which increases people's consumption budgets, results in lower utility growth than is implied by standard demand models that ignore peer effects.

These results also suggest that income or consumption taxes can have far lower negative effects on consumer welfare than are implied by standard models. This is because a tax that reduces my expenditures by 100 rupees will, if applied to everyone in my peer group, have the same effect on my utility as a tax of only 74 rupees that ignores the peer effects. This implies that about one-fourth of the money people might get back from an across the board tax cut does not increase utility, but instead is spent on increased perceived needs due to peer effects. The larger these peer effects are, the smaller are the welfare gains associated with tax cuts or mean income growth. We show this is particularly true to the extent that taxes are used to provide public goods or government services (that are less likely to induce peer effects themselves) rather than transfers.

We provide some calculations showing that the magnitudes of these peer effects on social welfare calculations, which are ignored by standard models of government tax and spending policies, can be very large. For example, we find potential welfare gains of hundreds of billions of rupees could be available in just a single existing government transfer program in India. We find similarly that the welfare gains in transfers from richer to poorer households (and more generally from progressive vs. flat taxes) may be much larger than previously thought, to the extent that poorer households do indeed have smaller peer effects than richer households.

REFERENCES

- Akay, Alpaslan and Peter Martinsson (2011), “Does relative income matter for the very poor? Evidence from rural Ethiopia.” *Economics Letters*, 110, 213–215. [1282]
- Armstrong, Timothy B., Michal Kolesár, and Mikkel Plagborg-Møller (2020), “Robust empirical Bayes confidence intervals.” arXiv preprint arXiv:2004.03448. [1279]
- Athey, Susan, Eckles Dean, and Imbens Guido (2018), “Exact p-values for network interference.” *Journal of the American Statistical Association*, 113, 230–240. [1278]
- Banerjee, Abhijit, Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson (2013), “The diffusion of microfinance.” *Science*, 341, 6144. [1258]
- Banks, James, Richard Blundell, and Arthur Lewbel (1997), “Quadratic Engel curves and consumer demand.” *Review of Economics and statistics*, 79, 527–539. [1264]
- Blackorby, Charles and David Donaldson (1994), “Measuring the cost of children: A theoretical framework.” *The Measurement of Household Welfare*, 51–69. [1263]
- Blume, Lawrence E., William A. Brock, Steven N. Durlauf, and Yannis M. Ioannides (2011), “Identification of social interactions.” In *Handbook of Social Economics*, Vol. 1, 853–964, North-Holland. [1260]
- Boneva, Teodora (2013), “Neighbourhood effects in consumption: Evidence from disaggregated consumption data.” Cambridge Working Papers in Economics 1328. [1258, 1262]
- Boskin, Michael J. and Eytan Sheshinski (1978), “Optimal redistributive taxation when individual welfare depends upon relative income.” *The Quarterly Journal of Economics*, 589 601. [1288]
- Bramoullé, Yann, Habiba Djebbari, and Bernard Fortin (2009), “Identification of peer effects through social networks.” *Journal of Econometrics*, 150 (1), 41–55. [1260]
- Brock, William A. and Steven N. Durlauf (2001), “Interactions-based models.” In *Handbook of Econometrics*, Vol. 5, 3297–3380, Elsevier. [1260, 1267]
- Carrell, Scott E., Richard L. Fullerton, and James E. West (2009), “Does your cohort matter? Measuring peer effects in college achievement.” *Journal of Labor Economics*, 27, 439–464. [1260]
- Chandrasekhar, Arun and Randall Lewis (2011), “Econometrics of sampled networks.” Unpublished manuscript, MIT [422]. [1259]
- Chao, Angela and Juliet B. Schor (1998), “Empirical tests of status consumption: Evidence from women’s cosmetics.” *Journal of Economic psychology*, 19, 107–131. [1258, 1262]
- Charles, Kerwin Kofi, Erik Hurst, and Nikolai Roussanov (2009), “Conspicuous consumption and race.” *The Quarterly Journal of Economics*, 124, 425–467. [1286]

- Clark, Andrew E., Paul Frijters, and Michael A. Shields (2008), "Relative income, happiness, and utility: An explanation for the Easterlin paradox and other puzzles." *Journal of Economic literature*, 46, 95–144. [1258, 1261]
- Clark, Andrew E. and Claudia Senik (2010), "Who compares to whom? The anatomy of income comparisons in Europe." *The Economic Journal*, 120 (544), 573–594. [1258, 1261]
- De Giorgi, Giacomo, Anders Frederiksen, and Luigi Pistaferri (2020), "Consumption network effects." *The Review of Economic Studies*, 87, 130–163. [1258, 1260, 1262]
- De Giorgi, Giacomo, Michele Pellizzari, and Silvia Redaelli (2010), "Identification of social interactions through partially overlapping peer groups." *American Economic Journal: Applied Economics*, 2, 241–275. [1260]
- Deaton, Angus (1988), "Quality, quantity, and spatial variation of price." *The American Economic Review*, 418–430. [1274]
- Diewert, Walter Erwin (1974), "Applications of duality theory." In *Frontiers of Quantitative Economics*, Vol. II (M. D. Intriligator and D. A. Kendrick, eds.), 106–171, North-Holland, Amsterdam. [1265]
- Donaldson, David and Krishna Pendakur (2006), "The identification of fixed costs from consumer behavior." *Journal of Business & Economic Statistics*, 24, 255–265. [1263]
- Duflo, Esther, Pascaline Dupas, and Michael Kremer (2011), "Peer effects, teacher incentives, and the impact of tracking: Evidence from a randomized evaluation in Kenya." *American Economic Review*, 101, 1739–1774. [1260]
- Dynan, Karen, and Ravina Enrichetta (2007), "Increasing income inequality, external habits, and self-reported happiness." *American Economic Review*, 97, 226–231. [1258, 1261]
- Easterlin, Richard A. (1974), "Does economic growth improve the human lot? Some empirical evidence." In *Nations and Households in Economic Growth*, 89–125, Academic Press. [1261, 1291]
- Frank, Robert. H. (1999), *Luxury Fever: Why Money Fails to Satisfy in an Era of Excess*. The Free Press, Simon and Schuster, New York. [1261]
- Frank, Robert. H. (2012), "The easterlin paradox revisited." *Emotion*, 12 (6), 1188–1191. [1261]
- Gali, Jordi (1994), "Keeping up with the Joneses: Consumption externalities, portfolio choice, and asset prices." *Journal of Money, Credit and Banking*, 26, 1–8. [1258]
- Geary, Roy C. (1949), "A note on a constant utility index of the cost of living." *Review of Economic Studies*, 18 (1), 65–66. [1262]
- Gorman, William M. (1976), "Tricks with utility functions." In *Essays in Economic Analysis: Proceedings of the 1975 AUTE Conference* (Michael J. Artis and Robert Nobay, eds.), 211–243, Cambridge University Press, Sheffield (Cambridge). [1263]

Gorman, William M. (1981), "Some engel curves." In *Essays in the Theory and Measurement of Consumer Behaviour in Honor of Sir Richard Stone* (Angus Deaton, ed.). Cambridge University Press, Cambridge. [1264]

Hansen, Lars. P. (1982), "Large sample properties of generalized method of moments estimators." *Econometrica*, 50 (4), 1029–1054. [1270]

Heffetz, Ori (2011), "A test of conspicuous consumption: Visibility and income elasticities." *Review of Economics and Statistics*, 93 (4), 1101–1117. [1286]

Jochmans, Koen and Martin Weidner (2016), "Fixed-effect regressions on network data." arXiv preprint arXiv:1608.01532. [1260]

Kahneman, Daniel (1992), "Reference points, anchors, norms, and mixed feelings." *Organizational behavior and human decision processes*, 51 (2), 296–312. [1257, 1261]

Kahneman, Daniel and Amos Tversky (1979), "Prospect theory: An analysis of decision under risk." *Econometrica*, 47 (2), 263–291. [1261]

Kalyanaram, Gurumurthy and Russell S. Winer (1995), "Empirical generalizations from reference price research." *Marketing Science*, 14 (3 supplement), G161–G169. [1258]

Lee, Lung-Fei (2007), "Identification and estimation of econometric models with group interactions, contextual factors and fixed effects." *Journal of Econometrics*, 140 (2), 333–374. [1258, 1260]

Lewbel, Arthur (1989), "A demand system rank theorem." *Econometrica*, 57 (3). [1264]

Lewbel, Arthur (1991), "The rank of demand systems: theory and nonparametric estimation." *Econometrica*, 711–730. [1264]

Lewbel, Arthur (1997), "Consumer demand systems and household equivalence scales." *Handbook of applied econometrics*, 2, 167–201. [1265]

Lewbel, Arthur, Samuel Norris, Krishna Pendakur, and Xi Qu (2022), "Supplement to 'Consumption peer effects and utility needs in India.'" *Quantitative Economics Supplemental Material*, 13, <https://doi.org/10.3982/QE1760>. [1265]

Luttmer, Erzo (2005), "Neighbors as negatives: Relative earnings and well-being." *Quarterly Journal of Economics*, 120 (3), 963–1002. [1257, 1258, 1261, 1262, 1263, 1288]

Manski, Charles F. (1993), "Identification of endogenous social effects: The reflection problem." *Review of Economic Studies*, 60 (3), 531–542. [1260, 1266, 1267, 1268]

Manski, Charles F. (2000), "Economic analysis of social interactions." *Journal of Economic Perspectives*, 14 (3), 115–136. [1260, 1267]

Ministry of Consumer Affairs, Department of Food and Public Distribution, Government of India (2018), "Salient features of the national food security act." Available at <http://dfpd.nic.in/Salient-features-National-Food-Security-Act.htm>. [1290]

NSS (2019), "Urban frame survey." Available at <http://mospi.nic.in/urban-frame-surveyuufs>. [1272]

Pendakur, Krishna (2005), “Semiparametric estimation of lifetime equivalence scales.” *Journal of Applied Econometrics*, 20 (4), 487–507. [1264]

Pollak, Robert A. and Terence J. Wales (1978), “Estimation of complete demand systems from household budget data: The linear and quadratic expenditure systems.” *American Economic Review*, 68 (3), 348–359. [1265]

Pollak, Robert A. and Terence J. Wales (1981), “Demographic variables in demand analysis.” *Econometrica*, 49 (6), 1533–1551. [1264]

Puri, Raghav (2017), “India’s national food security act (NFSA): Early experiences.” LANSAs, <https://assets.publishing.service.gov.uk/media/5964831e40f0b60a44000154/NFSA-LWP.pdf>. [1290]

Rabin, Matthew (1998), “Psychology and economics.” *Journal of Economic Literature*, 36 (1), 11–46. [1258]

Roth, Christopher P. (2014), “Conspicuous consumption and peer effects among the poor: Evidence from a field experiment.” (No. WPS/2014-29). [1286]

Roy, René (1947), “La Distribution du Revenu Entre Les Divers Biens.” *Econometrica*, 15 (3), 205–225. [1263]

Samuelson, Paul A. (1947), “Some implications of linearity.” *Review of Economic Studies*, 15 (2), 88–90. [1260, 1262]

Stevenson, Betsey and Justin Wolfers (2008), “Economic growth and subjective well-being: Reassessing the Easterlin paradox.” No. w14282, National Bureau of Economic Research. [1261]

Stone, Richard (1954), “Linear expenditure system and demand analysis: An application to the pattern of British demand.” *Economic Journal*, 64 (255), 511–527. [1262]

Veblen, Thorstein B. (1899), *The Theory of the Leisure Class, an Economic Study of Institutions*. Macmillan Publishers, London. [1261, 1286]

Co-editor Christopher Taber handled this manuscript.

Manuscript received 31 October, 2020; final version accepted 5 October, 2021; available online 19 January, 2022.