

Discussant comments on “What Do We Learn from Cross-Sectional Empirical Estimates in Macroeconomics” by Adam Guren, Alisdair McKay, Emi Nakamura, Jón Steinsson

Comments by Valerie A. Ramey, Revised May 13, 2020

1. Introduction

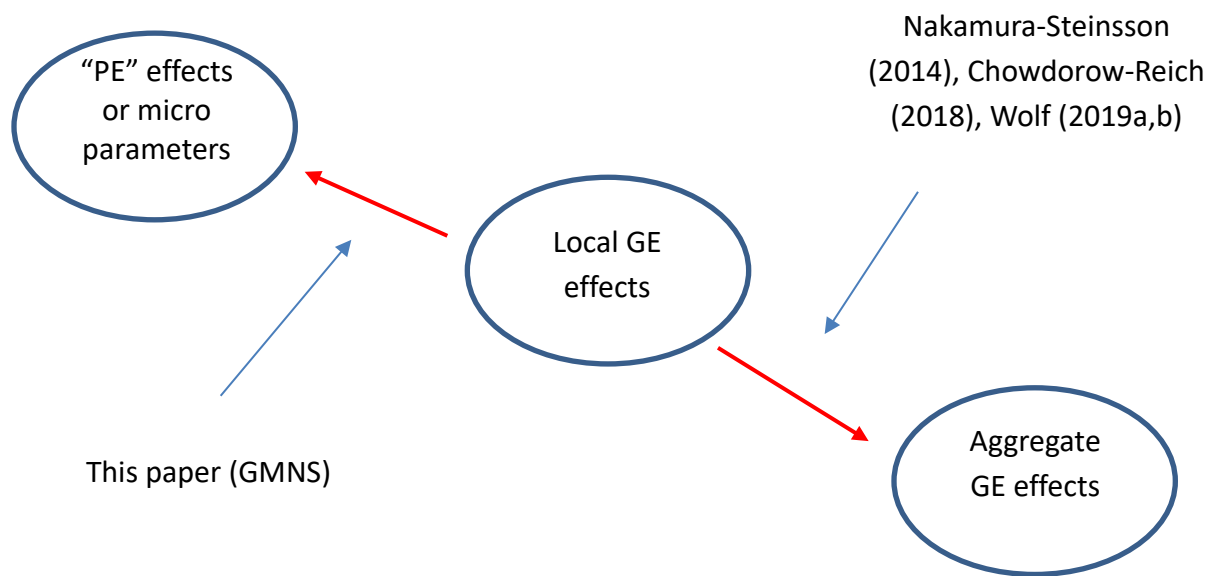
This insightful paper by Guren, McKay, Nakamura and Steinsson (GMNS) contributes to the literature that seeks to translate effects estimated on data at one level of geographic aggregation to another level of geographic aggregation. It can also be seen as a companion paper to their paper on the effects of variations of housing wealth on consumption (Guren, McKay, Nakamura, and Steinsson (forthcoming)). The current paper develops a very clever method for recovering “partial equilibrium” effects from regressions using variation across subnational units, such as cities and states. It then applies this method to estimating the marginal propensity to consume out of housing wealth. Separately, it offers a solution to a puzzle that has arisen with respect to the widely-used Saiz (2010) instrument used for house prices.

2. Linking Estimates at Different Levels of Aggregation

Identifying causal effects is particularly difficult in macroeconomics because of general equilibrium, dynamics, and expectations. Macroeconomists were slow to adopt the new methods developed in the applied microeconomics during its “credibility revolution.” The diffusion of applied micro methods to macro accelerated, however, at the start of the Great Recession, when researchers interested in macroeconomic questions began to estimate causal effects of fiscal policy, house price changes, etc. using cross-state or cross-city data. Initially, many thought that one could simply apply microeconomic methods to answer macro questions. The fiscal literature quickly realized, however, that the intercepts in cross-state (or city) regressions or time-fixed effects in state panel regressions netted out the macroeconomic effects we were trying to estimate.

This realization has led me to conclude that there is no “applied micro free lunch” for macroeconomists. Identification of macroeconomic effects can only be accomplished with macroeconomic identifying assumptions. As I outlined in Ramey (2019), there are two broad categories of methods for identifying macro causal effects: (i) using aggregate data in conjunction with time series identification or DSGE identification to estimate macroeconomic effects directly; or (ii) estimating micro or subnational effects and translating them to macroeconomic effects using macroeconomic theory, such as a DSGE model. GMNS’s method is related to the latter, but with an important twist.

The following diagram represents the effects at different levels of aggregation and how GMNS’s exercise differs from the standard exercise.



The rest of the literature has focused on determining the relationship between causal effects estimated at the state or local level (what GMNS call “Local general equilibrium effects”) to the aggregate general equilibrium effects. That is, they seek to translate the effects estimated for the middle circle to the downstream effects at the aggregate. GMNS instead seek to translate the local GE effects upstream to what they call “partial equilibrium” (PE) effects. As they make clear, what they are calling the “PE” effect is the effect of housing prices on consumption, holding other variables constant.

I think that the phrase “partial equilibrium” is misused here, as well as in numerous other places in the recent macroeconomics and public finance literature. Partial equilibrium refers to the *equilibrium* in one market, taking as exogenous prices in other markets as well as agents’ incomes. The *partial equilibrium* effect of an exogenous shock is the change in *equilibrium* price and quantity in that market. In contrast, the response of household or city-level average household consumption to a change in house prices, *ceteris paribus*, represents the *optimal responses* of individual households. It is a parameter of a best response function and it is based on the outcome of a *constrained maximization problem* for households. It is not a partial equilibrium outcome. Therefore, I will refer to the parameters they are trying to identify as *micro parameters*.

3. GMNS’s Method for Identifying Micro Parameters

GMNS explain their idea very clearly, first in a simple model and then in progressively more general models. In order to illustrate a later point I wish to make, I will explain their idea in a simple introductory macroeconomics Keynesian cross model. In particular, suppose the economy takes the following very simple form:

$$(1) Y = C + G \quad \text{NIPA identity}$$

$$(2) C = \alpha \cdot Y + \beta \cdot H \quad \text{Consumption function}$$

where Y is output (and income), C is consumption, G is government purchases, H is housing wealth, α is the marginal propensity to consume (MPC) out of income and β is the marginal propensity to consume out of housing wealth (MPCH).

Suppose we observe an exogenous increase in housing wealth in a city. How will consumption change? GMNS’s insight is that estimates across cities do not reveal the micro (what they call “partial equilibrium”) parameter β because the response estimated across cities contain endogenous feedback. To demonstrate their point, I solve the simple Keynesian cross (which assumes that output is demand-determined) to find equilibrium income and consumption in the city:

$$(3) \quad Y = \frac{\beta \cdot H + G}{1 - \alpha} = \frac{1}{1 - \alpha} [\beta \cdot H + G] \quad \text{City equilibrium output/income}$$

$$(4) \quad C = \frac{\beta}{1 - \alpha} H + \frac{\alpha}{1 - \alpha} G \quad \text{City equilibrium consumption}$$

These equations make clear that looking at the effects of exogenous changes in housing wealth across cities yields the following estimates:

$$(5) \quad \frac{dY}{dH} = \frac{dC}{dH} = \frac{\beta}{1 - \alpha}$$

As GMNS point out, the city consumption response to an exogenous increase in housing wealth confounds two effects: the marginal propensity to consume out of housing wealth, β , and the Keynesian multiplier, $1/(1-\alpha)$, which captures the local spillovers of one household's consumption on other households' income and consumption.

How then do we identify the marginal propensity to consumer out of housing wealth (MPCH) from estimates that use city variation? GMNS's insight is to identify the MPCH using an estimate of the government spending multiplier, i.e.

$$(6) \quad \frac{dY}{dG} = \frac{1}{1 - \alpha}$$

Thus, we can identify the MPCH (β) if we take a second step and divide the city estimate of the effect of an exogenous increase in housing wealth (from equation (5)) by an estimate of the government spending multiplier in equation (6).

GMNS's idea to use fiscal multipliers to help identify other key parameters is very clever. The idea is related to some independent contemporaneous work (e.g. Auclert and Rognlie (2020), Wolf (2019a,b)). GMNS generalize the idea with both analytic results and a dynamic macro model of multiple regions with potentially incomplete markets. Their impressive analysis shows that it is a relatively robust approximation under numerous generalizations.

4. Alternative Ways to Estimate Micro Parameters

It is useful to step back and think about why macroeconomists might want to estimate the marginal propensity to consume out of housing wealth, which I have argued is a *micro* parameter. After all, the title of the paper is about what we learn in macroeconomics. I agree with them that knowing the marginal propensity to consume out of housing wealth is useful for macroeconomists: it may be a key micro parameter in DSGE models that estimate the aggregate effects of housing booms and busts.

However, I would suggest that there are more straightforward ways to estimate this key micro parameter. This statement ties to my point in the earlier section where I argued that GMNS mischaracterize this parameter as a partial equilibrium effect when in fact it is a micro parameter. I will argue that while their method can be valuable in cases where there is no good micro data, in most cases there are better ways to estimate these types of parameters. These alternatives exploit rich household data and do not require a host of auxiliary assumptions. To demonstrate my point, consider two fine studies of the effect of changes in household wealth on household consumption.

The first study is by Campbell and Cocco (2007) who use UK household-level data and create a synthetic panel. They estimate a regression of changes in consumption on changes in house prices, controlling for household income, leverage and other demographic variables. Thus, their regression is analogous to GMNS's equation (2) and their estimate of the effect of housing wealth on consumption is exactly the micro marginal propensity to consume out of housing wealth, GMNS's parameter C_p . Campbell and Cocco's estimates are not confounded by endogenous changes in income in a city because their estimate is from a household-level regression that nets out city effects with a constant term and that furthermore controls for the household's income. Thus, the error term in the regression is entirely due to idiosyncratic factors of the household, not the endogenous city-level feedbacks that confound GMNS's estimates. Campbell and Cocco find that that elasticity of nondurable consumption to housing wealth varies from 0 for renters (as one would expect) to 1.7 for older homeowners. The average elasticity is 1.2, which roughly translates to a marginal propensity to consume out of housing wealth of 0.077.

The second excellent household-level study is by Aladangady (2017), who uses rich CEX data linked to confidential geographic detail in order to link household-level consumption to changes in household wealth in the U.S. His regression also controls for household income and other characteristics. He estimates a marginal propensity to consume out of housing wealth of 0 for renters and of 0.047 for homeowners.

But could household-level income be endogenous as well due to the local general equilibrium effect of house prices on income? In theory this is possible, but in practice there is no bias in most applications. There are two factors that eliminate the bias. First, the idiosyncratic components of household income swamp any common city components. For example, I estimate that city-fixed effects explain only 0.4 percent of the variation in household income in data from the 2000 Census. Second, the low estimates of marginal propensities to consume out of housing wealth mean that exogenous house prices are a minor source of variation in city-level income. My Monte Carlo simulations based on the simple Keynesian Cross model verify this intuition. I find that while estimates of the MPCH based on a city-level regression of consumption on exogenous house price are biased upward (by approximately the assumed government spending multiplier of 1.5), the estimates of the MPCH based on a household-level regression of consumption on both exogenous city-level house prices and household income show no bias.

Thus, data-rich and well-implemented household-level estimates solve the problem of the feedback from consumption to income in local general equilibrium that confounds cross-city estimates. The household-level approach obtains the estimates in one step. In contrast, consider how GMNS obtain the estimates in this paper and their companion paper. First, they estimate an elasticity of consumption to house prices of 0.072, based on a log difference specification using annual data on a panel of cities, instrument for housing prices. As they argue, this estimate includes the local GE effects. Thus, they must move to Step 2, which is to convert the elasticity to a local GE MPCH = 0.033 by dividing by average H/C ratio from 1985 to 2016, which was 2.17. Finally, in step 3 they purge this estimate of local GE effects by dividing by Nakamura-Steinsson's (2014) estimated state-level government spending multiplier of 1.5. Their final answer is a MPCH = 0.022.

But the review of two leading household level studies highlights a further challenge to GMNS’s approach. The micro evidence shows us that there is important heterogeneity in the MPCH. For obvious reasons, it is much higher for homeowners than renters. Thus, GMNS’s estimate of MPCH purged of local GE effects confounds the household-level response with the fraction of renters versus homeowners in the average city, i.e.,

$$MPCH_{city} = \lambda MPCH_{homeowners} + (1 - \lambda) MPCH_{renters}$$

where λ is the fraction of households that are homeowners. Because their baseline estimates do not weight cities by population, the λ implicit in their MPCH estimate is not necessarily nationally representative. Furthermore, it seems that calibrated DSGE models of housing would need separate estimates of the MPCH conditional on owning a home and the fraction of renters vs. homeowners in the economy.

5. Other Challenges to Implementing GMNS’s Method

GMNS consider many generalizations of their model. They are very upfront in showing that the approximation errors can be large for reasonable generalizations: the relative errors are often above 30 percent. In addition to the confounding of household MPCH’s with the fraction of renters that I discussed in the last section, there are other challenges to implementing their method. The main one is the reliance on the estimates of government spending multipliers.

As I discussed in an earlier section dY/dG , the government spending multiplier, is the linchpin estimate for their approach. It is the government spending multiplier estimate that allows conversion of a local GE estimate to a micro parameter of use for a macro model. For their illustration of their method, GMNS use estimates of the multiplier of 1.5 from Nakamura-Steinsson (2014), abbreviated “NS” hereafter.

When NS was published, the techniques used were state-of-the-art. However, technological progress has been so rapid in this literature in the last few years that we have figured out some of the things we did just a few years ago can be improved upon.

They study the effects of defense contracts on state (or region) level output and employment in a panel of states, annually from 1966-2006. To do this, they estimate the following regression:

$$\frac{Y_{it} - Y_{it-2}}{Y_{it-2}} = \alpha_i + \gamma_t + \beta \frac{G_{it} - G_{it-2}}{Y_{it-2}} + \varepsilon_{it}$$

where Y is per capita output in state i, G is per capita military procurement spending in state i. State and time fixed effects are included. β is the estimate of the multiplier. They consider two possible instruments: (i) an interaction of state dummies with aggregate procurement growth; and (ii) a Bartik instrument that interacts state historical sensitivity to aggregate procurement spending with the aggregate change in procurement spending.

When they use the interaction instruments (which are their preferred), they obtain a multiplier estimate of 1.43 (s.e. 0.36). However, these instruments have a low first-stage F-statistic, suggesting weak instruments. When they use the Bartik instrument, the multiplier is estimated to be 2.48 (s.e. 0.94), and the first-stage F-statistic is high. Substituting a multiplier of 2.5 rather than the 1.5 GMNS use has a significant effect on the implied marginal propensity to consume out of housing wealth (MPCH).

But there are other issues with the 1.5 multiplier estimate they import. First, the multiplier is not aligned on timing. Their city-level effects of house prices on consumption estimates are based on annual changes but the government spending multiplier estimates are based on biennial changes. To determine the impact of this timing, I used the Nakamura and Steinsson (2014) replication files to re-estimate their model using annual changes rather than their biennial changes. When I do this, the interactive instruments multiplier estimate falls from 1.43 to 0.69 and the Bartik instruments multiplier estimate falls from 2.48 to 1.65. Thus, aligning the timing means that one should use a smaller multiplier estimate to back out the MPCH.


I discovered additional issues with NS's estimates in the course of re-estimating their model. First, I discovered that the instruments are serially correlated – the correlation of the instrument (which is a two-year difference) and the previous two-year difference is 0.28. Because NS 2014 didn't include lagged Y's, G's, and instruments as controls, it means that their

estimates probably do not satisfy Stock and Watson’s (EJ 2018) lead-lag exogeneity condition. Vince Chen (2019) recognized this problem and tested for lead-lag exogeneity and rejected it for some states in the Nakamura-Steinsson data.

The best way to account for dynamics is to use the external instrument in an SVAR or in an IV-local projection (i.e. Ramey-Zubairy (2018)). Thus, I use the Nakamura-Steinsson data to estimate local projections using annual growth rates in order to align with the city data. I include two annual lags of the growth of output, government spending, and the Bartik instrument. It is important to include these additional lags. Vince Chen (2019) also estimated local projections using two-year growth rates, but did not include lags of the endogenous variables and instrument in his specification. Thus, his instrument is not a shock because it is forecastable from past values. See Stock and Watson (2018) and Alloza, Gonzalo, and Sanz (2019) for more discussion.

To estimate the integral multiplier at horizon h , I use the easy-to-implement one-step IV local projection (see, for example, Ramey-Zubairy (2018)). This integral multiplier is the ratio of the integral of the output response up to horizon h to the integral of the government spending response up to horizon h . Figure 1 shows the multiplier at each horizon, along with 95 percent confidence intervals (which correct for heteroskedasticity and autocorrelation).

Figure 1: Integral Multipliers by Horizon, Nakamura-Steinsson Data, IV-local projection



Several points are noteworthy about these estimates. The point estimates are 0.7 at one year, 0.75 at two years, and 1 at three years. Recall that the Nakamura-Steinsson static two-year multiplier estimate using the Bartik instrument was 2.48. Thus, using the same instrument but modeling the dynamics results in a fall in the multiplier from 2.48 to 0.75. Second, the multiplier appears to become larger and larger as the horizon grows. My analysis of the underlying impulse responses shows that it is the output response that keeps growing. It is not clear what the source of the rise is.

These new estimates do not deal with another potential issue that might affect the multiplier estimate needed to implement GMNS's procedure: the misalignment of geography. These government multiplier estimates are at the state level, but their housing wealth estimates are at the city level. Spillovers are more likely to be netted out by the intercept in city regressions than at higher level aggregation regressions, so this geographic misalignment could affect the procedure.

In sum, the GMNS method relies crucially on an estimate of the government spending multiplier. The point of this section is to demonstrate that those estimates can vary quite a bit, not only with the instruments used, but also with how dynamics are modeled, the horizon for the multiplier, and potentially with the geographic level.

6. Conclusions

Guren, McKay, Nakamura and Steinsson have introduced a useful new tool for identifying microeconomic parameters from cross-city or cross-state local general equilibrium estimates. While they illustrate their method in the context of estimating the micro-level marginal propensity to consume out of health wealth, it can potentially be used to identify micro parameters in a wide array of applications. Even when rich data makes household-level estimation the superior method, it can serve as a useful check.

As I have illustrated in my discussion, although GMNS have offered a great new "recipe," the quality of the final product is only as good as the ingredients. The two key ingredients of this

recipe are local general equilibrium effects estimates and local government spending multiplier estimates. Thus, high quality estimates of these two ingredients are crucial.

Supplemental References not included in GMNS References

Alloza, Mario, Jesús Gonzalo, and C. Sanz, "Dynamic Effects of Persistent Shocks," Banco de España Working Paper #1944.

Chen, Vince, "Fiscal Multipliers and Regional Reallocation," Boston University working paper, 2019.

Guren, Adam, Alisdair McKay, Emi Nakamura, and Jón Steinsson, "Housing Wealth Effects: The Long View," forthcoming *Review of Economic Studies*.

Ramey, Valerie A. "Ten Years after the Financial Crisis: What Have We Learned from the Renaissance in Fiscal Research?" *Journal of Economic Perspectives* 33, no. 2 (2019): 89-114.

Ramey, Valerie A., and Sarah Zubairy. "Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data." *Journal of Political Economy* 126, no. 2 (2018): 850-901.

Stock, James H., and Mark W. Watson. "Identification and Estimation of Dynamic Causal Effects in Macroeconomics using External Instruments." *The Economic Journal* 128, no. 610 (2018): 917-948.